

A CONCEPTUAL RESEARCH ESSAY

LOCATE YOURSELF

Platform Awareness

Know where you are before deciding where to go.

Davis Adedayo Eisape

Berlin · September 2026

platformgeneration.com

PUBLICATION INFORMATION

Locate Yourself

Platform Awareness

AUTHOR	Davis Adedayo Eisape
PUBLISHER	Platform Generation, Berlin, Germany
SERIES	Platform Generation Research Series, No. 01 / 2026
PUBLICATION DATE	September 2026
VERSION	1.0
WEB	platformgeneration.com daviseisape.com

© 2026 Davis Adedayo Eisape. Published by Platform Generation.
All rights reserved.
Short excerpts may be quoted for scholarly, journalistic and educational purposes with full attribution. All third-party images, archival documents and other reproduced materials remain subject to the rights and reuse conditions of their respective rights holders. Each figure is credited individually.

This essay is intended for research, reflection and strategic discussion. It does not constitute legal, financial or investment advice.

SUGGESTED CITATION

Eisape, D. A. (2026) *Locate Yourself: Platform Awareness*. Platform Generation Research Series, No. 01 / 2026. Berlin: Platform Generation.

ABSTRACT

Platform Awareness

Research on platforms has developed sophisticated accounts of network effects, two-sided markets, platform architecture, ecosystem governance, multi-platform integration and inter-platform ecosystems. Strategic analysis, however, often begins only after a focal platform has been selected and the organization's role relative to it has largely been assumed. This conceptual essay addresses the prior problem of orientation and proposes *Platform Awareness* as an organization's explicit and revisable understanding of the platforms in which it operates, the layer, position and role it occupies within each, and the strategic possibilities, dependencies and implications arising from those positions and their interconnection. The argument brings platform economics, platform and ecosystem strategy, infrastructure studies and situation-awareness research into conversation with a prehistoric boundary case and five historically grounded organizational cases: Edward Lloyd, Western Union, Sears, Amazon Web Services and Shopify. The cases are illustrative rather than representative and do not establish a causal relationship between awareness and performance. Building on the Platform Business Model Canvas, the essay subsequently proposes a Platform Landscape composed of connected PBMC Core Frames as a way of externalizing and contesting an organization's simultaneous positions across several platforms. The organizational capacity to develop, maintain and revise such awareness is treated as a related but distinct empirical question. The resulting research agenda therefore treats construct validity, role coding, negative cases, longitudinal change and possible capability formation as open questions rather than settled claims.

KEYWORDS

Platform awareness · platform strategy · platform ecosystems · inter-platform ecosystems · Platform Business Model Canvas (PBMC) · Platform Landscape

CONTENTS

The essay at a glance

1	The Question Before Strategy	1
2	Prehistoric Traces: Before Platforms Had Owners	12
3	Edward Lloyd: The Coffee Was the Product	21
4	Western Union: The Wire Carried No Money	27
5	Richard Warren Sears: Building on Other People's Infrastructure	33
6	Amazon Web Services: The Platform Was Already Inside	40
7	Shopify: Even a Platform Can Become a Participant	48
8	Platform Awareness: Naming the Pattern	54
8.1	Reading the cases as connected positions	55
8.2	Roles are relational, and infrastructure can disappear from attention	62
8.3	What the historical cases can establish	63
9	Locating the Organization: Platform Landscape and PBMC	65
9.1	The PBMC Core Frame as a repeated coordinate system	66
9.2	A Platform Landscape as connected Core Frames	67
9.3	The Landscape and the full PBMC correct one another	69
10	From Awareness to Strategy	71
10.1	Position before ownership	72
10.2	Landscapes move	73
10.3	Limits and a Research Agenda	73
	Conceptual Glossary	76
	References	82
	Figure Credits and Rights	93

The Question Before Strategy

Where, exactly, are we?

In the spring of 1974, a young economist at Bell Laboratories in Murray Hill, New Jersey, published a paper on a problem that almost nobody at the time would have described as the precursor to a new economic logic. Jeffrey Rohlfs had completed his doctorate at MIT only five years earlier and was a member of the technical staff at Bell Labs, an environment in which economists, mathematicians and engineers worked on highly practical questions in telecommunications. Rohlfs was studying the introduction of new communication services when he encountered a peculiar problem: the usefulness of such a service depended not only on its technical quality or its price, but also on how many other people chose to participate.

His paper carried the sober title *A Theory of Interdependent Demand for a Communications Service*. Its central intuition was already striking. The value of a communication service rises as more participants join. A communications system is of little use to an individual if hardly anyone can be reached through it; as the network grows, however, the value enjoyed by existing participants grows as well. Rohlfs did not yet speak of *network externalities*. He used the language of external economies in consumption, but he was already modelling a problem that would later become one of the building blocks of platform economics: a new service may work technically and be attractively priced and still fail if it does not reach sufficient critical mass (Rohlfs, 1974).

His example was neither a social network nor a digital marketplace. It was a new video-communications service. That historical distance is revealing. More than three decades before Facebook would make the size of a social network itself part of the product's value, a Bell Labs economist was trying to explain how a communications service could get started at all when every potential user had an incentive to wait until enough others were already there. A footnote in the original paper describes the young Rohlfs almost incidentally: Amherst, 1965; MIT doctorate, 1969; then Member of the Technical Staff at Bell Labs. A theory later frequently cited as foundational to network economics did not begin as a general theory of platforms. It began with the practical problem of how to launch a communications network whose value depended on its own growth (Rohlfs, 1974).

Eleven years later, two economists working at Princeton picked up related questions and pushed them much further. Michael Katz had completed his doctorate at Oxford; Carl Shapiro had earned his PhD at MIT in 1981 and taught at Princeton during the 1980s. Neither yet belonged to the senior generation of the discipline, although both would later move between academia and senior roles in US competition and regulatory policy. Katz served as Chief Economist of the Federal Communications Commission and as Deputy Assistant Attorney General for Economic Analysis in the Antitrust Division of the US Department of Justice; Shapiro likewise held senior positions in the Antitrust Division and later in the Council of Economic Advisers (Katz & Shapiro, 1985; U.S. Department of Justice, 2024a,b).

Their 1985 *American Economic Review* article, *Network Externalities, Competition, and Compatibility*, generalized the problem that Rohlfs had encountered in communications. A telephone is nearly useless if nobody else owns one and becomes more valuable with every additional person who can be reached. In computing, a similar dependency appears in another form: a machine can become more valuable when more compatible software exists, while software developers in turn have greater reason to invest where many users already own a particular type of hardware or operating system. Katz and Shapiro formalized such network externalities and examined how installed base, expectations and, in particular, compatibility shape competition among technologies (Katz & Shapiro, 1985).

The acknowledgements to their paper reveal that the idea did not emerge in an isolated office. Katz and Shapiro thank participants in discussions at Princeton, Columbia and Queen's University, as well as at the Federal Trade Commission and MIT. Before publication, the argument had already travelled between university seminars and a competition authority, touching precisely the boundary between theory and regulation that would later occupy both researchers professionally (Katz & Shapiro, 1985).

Their contribution was not to be the first to notice that networks can benefit from scale. Rohlfs had already modelled part of that logic eleven years earlier. Katz and Shapiro, however, developed a much broader account of competition in markets where the value of a product cannot be understood independently of the decisions made by other users, producers or developers. Under such conditions, a firm does not simply sell an isolated good. It offers access to a system whose attractiveness is created partly beyond the firm's immediate control.

The implications of that argument did not go uncontested. Less than a decade later, one can follow a scientific dispute almost page by page through a single issue of the *Journal of Economic Perspectives*. Katz and Shapiro published *Systems Competition and Network Effects* in 1994; Stanley Besen and Joseph Farrell followed with a discussion of standardization strategies; immediately after came an article by S. J. Liebowitz and Stephen Margolis whose title made clear that the authors were not merely offering a friendly extension: *Network Externality: An Uncommon Tragedy* (Katz & Shapiro, 1994; Besen & Farrell, 1994; Liebowitz & Margolis, 1994).

Liebowitz and Margolis did not deny the existence of network effects. Their criticism was aimed at the increasingly common conclusion that such effects almost automatically implied lock-in, market failure or a need for regulatory intervention. Some apparent externalities, they argued, could be internalized through ownership, contracts, prices or other market mechanisms, while the empirical foundation for some of the more dramatic stories about inefficient technological lock-in was surprisingly thin (Liebowitz & Margolis, 1994). The debate exposed an important distinction that still accompanies platform economics today. Interdependence among economic actors is first of all a description. Whether that interdependence generates market power, market failure or a case for regulation is a separate and far from trivial question.

While economists were still debating how far the consequences of network effects extended, the technological environment outside universities was changing rapidly. Fifteen years after Katz and Shapiro's first paper, some of the world's most visible technology companies were investing substantial sums in products that they then gave away for free.

In December 2000, the international information-systems community gathered in Brisbane for the 21st International Conference on Information Systems. The dot-com boom had captured companies, capital markets and the public imagination. Across the software industry, meanwhile, business decisions were proliferating that looked peculiar through the lens of conventional product logic. Microsoft and Netscape distributed browsers for free. Adobe charged nothing for Acrobat Reader. RealNetworks and Microsoft gave away media players. Sun Microsystems had acquired StarDivision in order to make software freely available.

Geoffrey Parker and Marshall Van Alstyne made this puzzle their problem. Both were still relatively early in their academic careers. Parker joined Tulane University as an Assistant Professor in 1998 after graduate training at MIT; Van Alstyne was then at the University of Michigan. Two young scholars were watching an internet economy whose rules were by no means settled and asking whether giving away digital products was really as irrational as it first appeared (Parker & Van Alstyne, 2000; Tulane University, 2011).

Their conference paper was titled *Information Complements, Substitutes, and Strategic Product Design*. Behind it lay an almost provocatively simple question: how can a company increase its profits by charging zero for a product? (Parker & Van Alstyne, 2000).

The answer went beyond the network-effects logic that had dominated earlier work. When two different groups of users are connected through complementary products or services, it can be rational for a firm to subsidize one side heavily, or even serve it for free, because growth on that side increases the value available on the other. The free browser, in other words, was not necessarily a product without a business model. It could be one component of a larger business model. Parker and Van Alstyne explicitly described these relationships as two-sided market externalities (Parker & Van Alstyne, 2000).

In hindsight, the logic can seem almost obvious. In 2000, it was not. *Platform* had not yet become a universal management term, and the behaviour of many internet companies was more often treated as a peculiarity of the New Economy than as the expression of a general market structure. Parker and Van Alstyne later developed their early ideas further; both became internationally prominent platform scholars and, together with Sangeet Paul Choudary, published *Platform Revolution* in 2016 (Parker et al., 2016). Yet their starting point was not an established platform paradigm. It was an economic puzzle unfolding in front of them.

At almost the same time, two French economists approached a closely related problem from another direction. Bernard Caillaud and Bruno Jullien gave their first joint CEPR paper, published in 2001, an unusually vivid title: *Chicken and Egg: Competing Matchmakers*.

Unlike Parker and Van Alstyne, Caillaud and Jullien were not recently graduated junior scholars. Caillaud had completed his doctorate at MIT and worked in contract theory and industrial organization; Jullien was a graduate of the École Polytechnique, ENSAE and EHESS, had earned his doctorate at Harvard and was already directing GREMAQ in Toulouse. In 1998 he received the CNRS Bronze Medal (Caillaud, 2019; Jullien, 2025). Two established industrial economists were therefore examining emerging internet markets with the tools of their discipline and encountering a problem that did not fit a simple one-sided logic.

Their “matchmakers” bring together two groups whose participation depends on one another. Buyers want to use a market where many sellers are already present; sellers want to be where many buyers are waiting. Each side would prefer the other to arrive first. The intermediary therefore has to do more than generate demand. It has to set two mutually dependent groups in motion at the same time (Caillaud & Jullien, 2001).

In June 2001, Jullien presented the paper at a joint CSIO/IDEI workshop in Toulouse. On the same programme was a second presentation whose title, in retrospect, reads almost like the formal naming of a field: Jean-Charles Rochet presented *Platform Competition in Two-Sided Markets* (Toulouse School of Economics, 2001).

The fact that several researchers were recognizing closely related structures at almost the same moment makes the emergence of platform economics more interesting than a conventional invention story. Nobody sat down at a desk and invented the theory of platforms in a single act. Around the turn of the millennium, enough markets appeared to display the same unusual logic that scholars working independently began to see a common pattern.

Rochet and his co-author Jean Tirole were working in an especially fertile setting. Rochet had originally trained as a mathematician at the École Normale Supérieure before moving into banking, insurance and financial economics. Tirole had first trained as an engineer at the École Polytechnique and the École nationale des ponts et chaussées before completing a doctorate in economics at MIT. By the early 2000s he was already one of the world’s leading scholars in industrial organization; in 2014

he received the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel for his analysis of market power and regulation (Nobel Prize Outreach, 2014).

That one of the formative platform papers was co-authored by a future Nobel laureate is less revealing, however, than the fact that Rochet and Tirole were not initially searching for a general theory of digital platforms. They were studying credit cards.

Around the economist Jean-Jacques Laffont, GREMAQ and later the Institut d'Économie Industrielle (IDEI), Toulouse had developed an unusual research environment in which economic theory was deliberately connected to practical problems in telecommunications, banking, payments, energy and other regulated industries. Companies and institutions brought real problems into this environment; researchers attempted to identify more general economic structures within them. Industry problems did not merely appear as later applications of a completed theory. They could become the starting point from which new theory emerged (Toulouse School of Economics, 2018).

Two decades later, Rochet and Tirole described how this happened in their own work. In the payments industry, they had initially been concerned with how fees should be distributed between cardholders and merchants. A payment card becomes more attractive to consumers when many merchants accept it; acceptance becomes more attractive to merchants when many customers hold the card. While analysing this pricing structure, the two economists realized that the same logic appeared in operating systems, internet platforms, media and video games (Rochet & Tirole, 2023).

The paper's development can be traced with unusual precision. Rochet presented *Platform Competition in Two-Sided Markets* in Toulouse in early June 2001; a surviving working-paper version by Rochet and Tirole is dated 26 November 2001, and the idea that a platform must bring different groups "on board" together was already central to the emerging argument. Rochet and Tirole then submitted the paper to a leading general-interest economics journal. According to their later recollection, the referee reports were very favourable, yet the editor rejected it because, as they recalled two decades later, "*the topic was too narrow*" for a generalist journal (Rochet & Tirole, 2023). The subsequent publication history gives that judgment an institutional irony. Tirole was President of the European Economic Association in 2001 and, together with other officers, had recently helped create a new journal for the association after dissatisfaction with the publishing arrangement surrounding the *European Economic Review*. Rochet and Tirole submitted the rejected paper to the new *Journal of the European Economic Association*; although Tirole belonged to the team behind the journal, the authors later explicitly noted that he did not handle his own submission. In 2003, *Platform Competition in Two-Sided Markets* appeared in Volume 1, Issue 4 of the newly founded *JEEA* and became one of the defining texts of the emerging field (Rochet & Tirole, 2003, 2023).

The paper's decisive theoretical move was not simply to place two groups on opposite sides of a diagram. Rochet and Tirole showed that in such markets not only

the overall level of prices matters, but also the distribution of prices across the sides. A platform may heavily subsidize one group, or serve it for free, while earning a substantial share of its revenues from another if that asymmetric price structure is what enables sufficient participation on both sides. In their later *Two-Sided Markets: A Progress Report*, they sharpened the formulation: a market is two-sided in this sense when not only the aggregate price level but the price structure affects the volume of transactions (Rochet & Tirole, 2006).

Unlike the earlier disputes over network externalities, there was no similarly prominent moment of fundamental rejection. More striking was the speed with which the perspective spread. David Evans applied it to antitrust questions in a 2002 working paper. That paper was initially titled *The Antitrust Economics of Two-Sided Markets*; by the time a version appeared in the *Yale Journal on Regulation* in 2003, the title had broadened to *The Antitrust Economics of Multi-Sided Platform Markets*. The vocabulary of the field was changing in real time: from two-sided markets toward multi-sided platforms (Evans, 2003).

Evans discussed dating clubs, credit cards, operating systems and other intermediaries as variants of the same economic problem, while warning against applying conventional antitrust benchmarks mechanically to markets in which one side may rationally be served at a price of zero, or even below its immediate cost. Mark Armstrong published another influential model of competition in two-sided markets in 2006; by 2009, Marc Rysman could review the subject in the *Journal of Economic Perspectives* as an identifiable and rapidly growing field of economics (Armstrong, 2006; Rysman, 2009).

By the time this vocabulary had begun to stabilise, firms built around related platform logics were also becoming increasingly prominent at the top of global equity markets. Figure 1.1 deliberately juxtaposes two different kinds of historical information without turning them into a single numerical series. The blue line is a narrow measurement of economic salience: it reports how much market value was represented by classified platform firms within the world's ten most valuable listed companies at each observation date, while the smaller percentage labels show those firms' share of total Top-10 market capitalisation. Beneath the quantitative axis, selected intellectual milestones extend the chronology much further back. Their purpose is not to imply that Rohlfs in 1974, Katz and Shapiro in 1985, or the other early authors were already studying "platforms" in today's fully developed sense. Rather, the juxtaposition makes visible a historical asymmetry that motivates this essay: mechanisms that later became foundational to platform economics were theoretically observable decades before platform firms became economically prominent among the world's largest listed companies. The figure therefore raises, rather than resolves, a question that will recur throughout the historical cases: how much of what now appears new is genuinely new, and how much became visible only after a vocabulary for recognising it had developed?

The theoretical vocabulary resonated because it made comparable a set of arrangements that had long been described separately. Credit cards predated the




Figure 1.1. Selected intellectual milestones and the measured economic salience of classified platform firms. The chronology deliberately begins decades before the market-capitalisation series: Rohlfs (1974) analysed interdependent demand in communications long before today’s platform firms had become economically prominent. The juxtaposition does not imply that Rohlfs studied platforms in the contemporary sense; it makes visible a central puzzle of this essay: mechanisms that later became foundational to platform economics were theoretically observable before the corporate forms now associated with them became economically dominant. The figure therefore raises rather than resolves the question of whether platform structures are genuinely new, or whether the vocabulary for recognising them emerged later than some of the structures themselves.

Note: The market-capitalisation series begins in 2009. The vertical axis reports, in trillions of US dollars, the combined market capitalisation of companies classified at that date as transaction, innovation, or hybrid platform firms using the distinction developed by Cusumano et al. (2019). The labels report the absolute value and the same firms’ share of the combined market capitalisation of the ten most valuable publicly listed companies globally on 31 March of the respective year. Historical market-capitalisation data are drawn from PwC’s *Global Top 100* series, based on Bloomberg data for 2009–2020 and S&P Global Market Intelligence for 2026 (PricewaterhouseCoopers (PwC), 2014, 2016, 2019, 2021, 2026). The blue line is therefore a measured indicator of economic salience among the largest listed firms, not an estimate of the total size of the platform economy and not evidence that platform status caused market-value growth. Historical classifications are date-specific rather than retrospectively applying a firm’s current architecture. The names beneath the zero line are selected intellectual milestones, not observations on the vertical scale and not a quantitative series of research activity; they mark publication years chosen to orient the intellectual history developed in this chapter. Their placement before 2009 is intentionally descriptive rather than causal: it shows that parts of the theoretical vocabulary and several mechanisms later associated with platform economics were being analysed before the measured corporate series begins.

commercial internet, newspapers had connected readers and advertisers for generations, shopping centres coordinated retailers and consumers, and video-game consoles brought software developers and players into mutually dependent markets. What changed around the turn of the millennium was not the sudden appearance of interdependence, but the growing ability to recognize apparently different industries as manifestations of related economic and organizational structures.

The resulting field nevertheless developed along several intellectual paths. Economists concentrated on pricing, network effects and market sides; information-systems scholars examined digital infrastructures, standards and technical architectures; strategy researchers studied innovation, industry structure and governance. Rietveld and Schilling later showed how far these vocabularies diverged even where the underlying questions overlapped: what one literature called multihoming could appear elsewhere as compatibility, exclusivity or standards competition (Rietveld & Schilling, 2021). Platform research therefore became interdisciplinary without immediately becoming integrated, and some of its most important advances came from scholars who crossed the boundaries between these traditions.

Annabelle Gawer's early work on Intel was one such bridge. When she submitted her MIT Sloan dissertation in 2000, her starting point was not the price structure between two market sides but the organizational problem faced by a technically powerful firm whose products depended on complementary innovation by companies outside its hierarchy. Intel could improve its processors and still be constrained if hardware manufacturers, software firms or other complementors failed to advance with it. Gawer's subsequent work with Michael Cusumano helped turn that observation into a broader account of platform leadership, in which the focal firm does not merely optimize a product but also shapes interfaces, incentives and conditions under which other actors innovate (Gawer, 2000; Gawer & Cusumano, 2002, 2014). The platform, in this tradition, was already difficult to understand by looking only inside the firm.

The parallel rise of ecosystem research widened the analytical boundary still further. Ron Adner, who had originally trained as a mechanical engineer, increasingly emphasized that a technically excellent innovation can fail because the surrounding activities and partners required for a value proposition are not ready when the focal firm is ready (Adner, 2012). By 2017 the term *ecosystem* had become so popular that its breadth itself created a conceptual problem. Adner described a "cornucopia of constructs" in which ecosystems, platforms, networks, supply chains and business models were frequently invoked without sufficient precision, and responded by proposing an ecosystem-as-structure approach organized around a focal value proposition, the activities required to realize it, the actors performing those activities, the positions they occupy and the links among them (Adner, 2017, p. 39). Jacobides, Cennamo and Gawer subsequently asked under what conditions ecosystems constitute a distinct organizational form and how legally independent

actors can be coordinated through complementarities, modularity and governance without being brought wholly inside a conventional hierarchy (Jacobides et al., 2018).

Once research had moved beyond the legal boundary of the focal firm, however, a second boundary became increasingly visible: the tendency to make one platform or one ecosystem the centre of the analysis. Petra Bosch-Sijtsema and Jan Bosch stated the limitation unusually directly in 2015. Studies, they observed, often discussed one network and selected members within it even though firms can work in several ecosystems simultaneously and may need different engagement models in each; their qualitative cases showed organizations taking different positions, and therefore possessing different degrees of influence, across those ecosystems (Bosch-Sijtsema & Bosch, 2015). Staykova and Damsgaard approached a related problem through what they called *platform constellations*, arguing that highly connected platforms governed by the same provider should not always be studied as independent entities (Staykova & Damsgaard, 2016). Tavalaei and Cennamo later examined the problem from the complementor side, showing that a participant may have to decide whether to specialize within one platform ecosystem or distribute its activities across several (Tavalaei & Cennamo, 2021).

The issue is broader than digital firms maintaining a portfolio of products. Shipilov and Gawer begin their review of interorganizational networks and ecosystems from the observation that organizations are simultaneously embedded in both, while the corresponding research traditions had developed to a considerable degree in isolation; their call is therefore not simply for more ecosystem research, but for analytical approaches capable of making patterns of interdependence across organizational structures more visible (Shipilov & Gawer, 2020). Schrieck and colleagues make an analogous point within the platform literature itself. In developing a typology of multi-platform integration strategies, they note that previous work had concentrated predominantly on the launch, growth and competition of single platforms, even as firms such as Alphabet, Apple, Meta, Alibaba and Tencent increasingly operated portfolios in which several platforms were deliberately integrated (Schrieck et al., 2024).

Research on inter-platform ecosystems takes another step because the platform itself is no longer treated only as the centre around which non-platform actors gather. Carballa-Smichowski, Duch-Brown, Martens and Gomez-Losada analyse 241 European platforms and describe configurations in which platforms become complementors to one another, direct users across platform boundaries, aggregate other platforms or host interactions that create externalities between separate systems. Importantly for the argument developed here, such configurations need not possess one overarching owner who orchestrates the entire arrangement (Carballa-Smichowski et al., 2026). The analytical world therefore becomes considerably less tidy than the familiar picture of one Platform Owner surrounded by its users and complementors.

These developments mean that the research gap cannot credibly be described as an absence of work on multiple platforms. Multihoming, participation in several ecosystems, platform constellations, multi-platform integration and inter-platform

ecosystems all address parts of that terrain. What remains less systematically developed is the prior problem faced by an organization whose strategic starting point is not yet a clearly specified platform: how can it identify the several platforms in which it is simultaneously situated, distinguish the different positions it occupies across them, and interpret the dependencies and possibilities created by their interconnection? Much platform research quite reasonably begins by selecting a focal platform and asking how that system works. The inquiry in this essay begins one step earlier, with the organization that is trying to understand which systems should be focal in the first place.

A nearby term already exists in the literature. Burström, Grönlund and Ranti use *situational platform awareness* in a managerial context concerned with understanding an identified platform, including its market role, components, interfaces and important interdependencies (Burström et al., 2016). The present argument begins one analytical step earlier and widens the object of orientation: which platforms should be treated as relevant in the first place, and what different position does the same organization occupy across them? *Platform Awareness* as used here is therefore conceptually related to, but not synonymous with, situational platform awareness.

The term *platform awareness* has also been developed independently in media-education research. Spurava uses it to describe a comprehensive understanding of the interplay between the social, technological, economic and political dimensions of platform power, with the purpose of strengthening digital media users' agency and related literacies (Spurava, 2024). That individual and media-literacy-oriented construct addresses a different analytical object from the organization-centred strategic orientation developed here. The overlap in terminology therefore does not imply conceptual equivalence: the present essay concerns an organization's explicit and revisable understanding of its positions across several platforms.

That shift of analytical anchor changes what counts as a useful description of the firm. A company may operate a technical platform internally while orchestrating a customer-facing ecosystem, consume infrastructure from another platform, provide capabilities into a fourth system and depend on standards or institutions whose strategic importance becomes visible only when they change. None of those positions defines the organization in isolation, because each exists relative to a particular platform and the value-generating relationships organized around it. The question whether a firm is an Owner, Provider, Consumer, Partner or, in the terminology of other platform literatures, a complementor is therefore incomplete until the reference system is specified.

This is also why platform strategy need not begin with the aspiration to build or own a platform. Ownership can confer valuable rights over interfaces, access and governance, but organizations can also create substantial value by occupying strong positions in systems controlled by others, by externalizing capabilities developed for internal use, or by recognizing that a new platform layer is forming around a business they already orchestrate. Before any of those choices can be evaluated, a

more elementary orientation problem has to be solved: which platforms already matter to the organization, and who is the organization within each of them?

In this essay, *Platform Awareness* is used provisionally for that orientation problem: an organization needs an explicit and revisable understanding of which platforms matter to it, where it is positioned within them and what follows from their interconnection. Chapter 8 formalizes the concept after the historical cases have made the recurring positional movement visible. Edward Lloyd did not need a canvas to understand that maritime information and commercially relevant people were gathering in his room, and Richard Sears did not need an ecosystem map before exploiting railway and postal infrastructures. The purpose of the later tools is not to create the phenomenon retrospectively, but to make an organization's understanding explicit enough to compare, question and revise.

Before developing that framework, however, it is useful to move much further back than Rohlfs, Katz, Shapiro or Tirole. If the platform perspective ultimately concerns structured interdependence among actors, resources, recurring interactions, interfaces and rules, then the history relevant to the argument does not begin with digital technology, and perhaps not even with markets in any recognizable modern sense.

A Note on Method

This is a conceptual essay rather than an empirical test of a finished construct. It develops the idea of Platform Awareness by bringing several established research streams into conversation first with a prehistoric boundary case and then with five historically grounded organizational cases. The prehistoric discussion is used to test how far a technology-neutral language of platform-like coordination can be carried without retrospectively declaring ancient social arrangements to be platforms. The five organizational cases were chosen because each exposes a different kind of positional shift: a surrounding system becomes strategically visible, an existing network enables a new transaction, infrastructure owned by others becomes part of a business model, an internal capability becomes an external platform, or a platform itself becomes a participant in a larger system. They are illustrative cases, not a representative sample.

The historical sections therefore do not try to prove that Platform Awareness caused success. Nor do they attribute thoughts or motives to historical actors unless a contemporary or retrospective source supports them. Where the evidence ends, the language changes with it: documented fact is separated from plausible inference, and both are separated from the platform interpretation offered here. The purpose is narrower, but useful. The cases give us concrete situations in which to ask whether an organization's strategic possibilities can look different once we change the level at which we locate it.

CHAPTER 2

Prehistoric Traces: Before Platforms Had Owners

OHALO II AND GÖBEKLI TEPE

Deep antecedents extending back into prehistory.



Figure 2.1. Ohalo II makes early cereal harvesting visible; Göbekli Tepe makes large-scale coordination and monumental construction visible.

PLATFORM LOGIC DOES NOT REQUIRE DIGITAL TECHNOLOGY

In the autumn of 1989, drought at the Sea of Galilee revealed something that had remained hidden for more than twenty millennia. Several years of low rainfall, intensified by heavy water extraction, had lowered the lake so far that stretches of its south-western shore normally submerged beneath the water emerged into the open. On one of these newly exposed beaches, a team led by Dani Nadel, then a relatively young prehistorian at the University of Haifa, encountered the traces of an Ice Age camp. Nadel had not yet completed his doctorate on spatial patterns at prehistoric sites. Ohalo II would become a place he returned to for decades and with which his name would later be closely associated (Nadel, 1990, 2017).

Archaeologically, the drought was a gift. Socially, of course, it was anything but. Even so, it is difficult not to imagine the strange ambivalence of later dry years. Farmers, communities and authorities hoped for rain; archaeologists around the lake must also have watched the water-level reports with particular attention. After the first excavation campaigns, Ohalo II disappeared again beneath the rising lake, and further work became possible only when the water level fell once more at the end of the 1990s. Not every subsequent drought produced a second Ohalo II moment. The discovery of 1989 was enough, however, to alter significantly what archaeologists thought they knew about the lives of Ice Age hunter-gatherers.

Even the site's name contains a small piece of that history. About 800 metres farther north, an earlier period of low water in 1986 had already exposed a prehistoric site rich in flint and basalt artefacts. It was later named Ohalo I. When a second site emerged in 1989, it became Ohalo II. Ohalo I turned out primarily to be a broad working area with traces from different phases of use. Ohalo II, by contrast, offered something much rarer because of its exceptional preservation: not only stone tools, but part of a human living environment (Nadel, 2006).

The same water that made discovery difficult was also responsible for that preservation. Fine sediments covered the site; oxygen-poor conditions slowed the decay of organic material. Beneath the mud lay the remains of six oval huts made from branches and other plant matter, hearths, stone tools, animal and fish bones, shell beads, a human grave and large quantities of charred seeds and fruits. One hut even contained grasses apparently used as floor covering or bedding, among the oldest known evidence of such sleeping material (Nadel et al., 2004).

Ohalo II is approximately 23,000 years old. The first communities in the Near East that were durably organized around farming emerged roughly 11,000 to 12,000 years ago. In other words, around eleven to twelve thousand years separate the people of Ohalo II from those early farmers. Roughly the same span separates those early Neolithic farming communities from us. Ohalo II therefore does not sit just before agriculture. It lies approximately one entire "Neolithic-to-present" interval earlier.

That chronological distance is precisely why some of the finds were so unsettling. Five flint blades, for example, hardly look like objects capable of shifting a long chronology of human technology. For Iris Groman-Yaroslavski, however, their most important surface existed at the microscopic scale: fine polish, patterns of abrasion and damage from which it is possible to reconstruct what a stone tool was actually

used to do. At the University of Haifa, the archaeologist specialized in precisely this form of use-wear analysis and later led the relevant laboratory at the Zinman Institute. Together with Nadel and the archaeobotanist Ehud Weiss, she examined the blades and concluded that they had been used to cut semi-ripe wild cereals shortly before the natural dispersal of the grains. Some blades had apparently been inserted into handles and formed part of composite harvesting tools (Groman-Yaroslavski et al., 2016).

Weiss arrived at the same site from a very different route. He had studied bioarchaeology and geoarchaeology at University College London, completed a doctorate in archaeobotany in Israel and, in the early 2000s, worked as a postdoctoral fellow in Ofer Bar-Yosef's Stone Age Laboratory at Harvard University. What lay before him there was not the usual modest collection of charred seeds that archaeobotanists patiently extract from sediments, but something close to an absurdly rich Ice Age plant archive: tens of thousands of remains that under normal conditions would have disappeared long ago. Weiss later recalled, in essence, that before Ohalo II he had scarcely imagined that the plant diet of such an ancient community could be reconstructed in that level of detail; unexpected finds of this kind, he told the *Harvard Gazette*, were part of the reason one does research in the first place (Weiss et al., 2004).

Groman-Yaroslavski could read from the tools how plants had been cut; Weiss could reconstruct from the plant remains which plants had been gathered, processed and eaten. Together they produced a picture far less primitive than a simple opposition between "forager" and "farmer" would suggest. The five blades are considered the earliest known evidence of composite cereal-harvesting tools, around 8,000 years before the Natufian period and about 12,000 years before the early farming communities of the Neolithic (Groman-Yaroslavski et al., 2016).

The people of Ohalo II were not farmers. They hunted, fished and used a remarkably broad range of wild plants. Yet those plant remains increasingly made the site uncomfortable for overly neat historical categories. Among roughly 150,000 specimens were wild wheat, wild barley and wild oats, together with seeds from plant species that today are especially common in human-disturbed ground and cultivated fields.

Ainit Snir, then a doctoral researcher working in Ehud Weiss's academic environment and later the author of a dissertation on the construction, use and preservation of the Ohalo II huts, placed these plants at the centre of a 2015 study with Nadel, Weiss and other researchers. They called the species proto-weeds: possible predecessors of the weed communities that would later become an almost unavoidable companion of cultivated fields (Snir et al., 2015).

It was an appealing observation because weeds are not usually among the things with which civilizations choose to decorate their histories. For archaeobotanists, however, they can be extraordinarily informative. Certain plants thrive particularly well where people break up soils, disturb vegetation and return repeatedly to the same areas. Together with harvesting tools, grinding stones and large quantities of wild cereals, Snir and her colleagues therefore interpreted the proto-weeds as a

possible indication that the people of Ohalo II were already experimenting, on a limited scale, with the cultivation of wild cereals (Snir et al., 2015).

The authors were more cautious than some later headlines suggested. They did not claim to have found domesticated wheat at Ohalo II, nor did they construct a direct line of descent from the people at the site to Neolithic farmers. Perhaps something was tried there and then disappeared again. Perhaps the boundary between intensive gathering, deliberate intervention in plant populations and cultivation was itself much less sharp than the categories we use retrospectively imply.

That possibility changes how innovation looks. Not every discovery has to be passed from generation to generation; not every experiment has to open a new historical era. People can recognize a possibility, use it and abandon it. The same possibility may become attractive again thousands of years later under different ecological or social conditions.

How much of later agriculture was deliberately planned and how much emerged through long-term feedback between human behaviour and ecological change remains one of the more interesting disputes in research on domestication. Robert Spengler, an American archaeobotanist who later studied the relationships among plants, trade and human migration at the Max Planck Institutes in Jena, for example emphasizes the selective effects of human behaviour itself. People who collect, transport, store and intentionally or unintentionally redistribute seeds are already altering the environment in which plants compete, even without pursuing a breeding programme that spans millennia (Spengler, 2020).

Shahal Abbo approaches the process from a somewhat different direction. Abbo is not only an archaeobotanist but also a plant scientist and geneticist at the Hebrew University of Jerusalem who has worked extensively on the actual domestication and breeding of crops, including chickpeas. Together with the archaeologist Avi Gopher, he warns against models of unintended evolution that reduce early cultivators almost to bystanders in their own history. People who select seeds over years, observe sowing times and tend plant populations are likely to possess considerable practical knowledge. The transition to agriculture, in their view, cannot be explained solely as an ecological side effect of human activity (Abbo & Gopher, 2020).

The two positions do not have to be mutually exclusive. Humans act intentionally without being able to foresee every consequence of their actions. They alter their environment, and the altered environment then creates new possibilities and constraints. This feedback leads to a concept that would later become influential well beyond evolutionary biology: niche construction.

The term was coined in 1988 by the British biologist John Odling-Smee, then working in Oxford. Odling-Smee was dissatisfied with an image of evolution in which organisms largely respond to an already given environment. Beavers build dams, earthworms alter soils, humans clear forests and irrigate fields. In each case, organisms themselves alter environmental conditions that subsequently affect natural selection. With collaborators, he developed the argument through the 1990s and early 2000s into a broader theory and elaborated the idea of ecological inheritance:

descendants inherit not only genes, but in part a world already altered by earlier generations (Odling-Smee, 1988; Odling-Smee et al., 2003).

Odling-Smee remained associated with this idea for decades and later worked as an Emeritus Research Fellow in the orbit of Mansfield College and the University of Oxford. The subtitle of his 2003 book with Kevin Laland and Marcus Feldman, *The Neglected Process in Evolution*, was programmatic. The authors believed that an important part of evolutionary causation was visible everywhere and yet theoretically underestimated.

Not every biologist welcomed the suggestion that a “neglected process” had been discovered. Critics pointed out that Darwin himself had written about earthworms and that ecologists had long known organisms alter their environments. The dispute was therefore less about whether organisms change environments than about whether those changes deserved a distinct theoretical status within evolutionary biology. There is an instructive parallel here with later platform research: sometimes the innovation lies not in discovering a completely unknown phenomenon, but in recognizing scattered observations as manifestations of the same structure.

In humans, such feedback becomes especially far-reaching because we transmit not only altered physical environments but also tools, knowledge, norms and institutions. A generation therefore inherits more than genes and landscapes. It inherits a world in which others have already laid paths, spread plants, created rules and established expectations.

This is where Yuval Noah Harari enters the story, although less as an archaeologist than as an exceptionally successful narrator of very large historical patterns. Harari, born in Israel in 1976, completed his doctorate at Oxford in 2002 not on the Stone Age or agriculture but on medieval military history, initially specializing in the culture of warfare in the Middle Ages and early modern period. Only later did he begin teaching a world-history course at the Hebrew University of Jerusalem. Those lectures eventually became the Hebrew book known internationally as *Sapiens*. A medieval historian thus became, almost incidentally, one of the world’s most widely read writers on the entire history of *Homo sapiens*.

That success gave Harari millions of readers and an unusual position between scholarship and public intellectual life. Specialists have praised his ability to make large patterns visible while repeatedly criticizing the way some of his strongest sentences can sound more certain than the research on which they rest. For that reason, he is useful here less as an evidentiary authority than as someone capable of sharpening a perspective until it becomes hard to forget.

One of his best-known reversals concerns wheat. The conventional story says that humans domesticated grain. Harari asks us to turn the perspective around. In his deliberately provocative formulation: “*We did not domesticate wheat. It domesticated us.*” (Harari, 2015, p. 81). Wheat, of course, has no intentions. Yet people who became increasingly dependent on it had to work soils, remove weeds, provide water, protect crops and organize more of their everyday lives around the needs of a plant that had previously been only one among many food sources.

As historical rhetoric, the reversal is stronger than as archaeological theory. Yet it exposes a structure. Humans change a system, create new opportunities within it, and then begin to reorganize their own behaviour around the system they have changed.

Ohalo II does not show this process as a farming society. That is exactly why the site is so interesting. There is no clean threshold at which a forager goes to sleep one evening and wakes the next morning as a farmer. Instead we see techniques, environmental interventions, experiments and habits, some of which later become central and others of which may disappear again.

Even the human skeleton from Ohalo II warns against confusing a good story too quickly with a secure one. In a shallow grave lay the body of a man approximately 35 to 40 years old. An early osteological analysis interpreted the pronounced asymmetry of his arms as possible paralysis, and therefore as evidence that his community may have supported a physically impaired person over a substantial period. Later comparisons questioned the diagnosis itself and argued that the asymmetry might still fall within the range of variation seen in Palaeolithic skeletons (Hershkovitz et al., 1995).

The grave itself did not change, but the scientific story told about it did, a useful reminder that interpretation can move even when the material evidence remains fixed. Almost twelve thousand years after Ohalo II, several hundred kilometres to the north, Göbekli Tepe would make the same problem of changing interpretation visible on a much larger archaeological scale.

The first scientific encounter with the hill was anything but spectacular. In 1963, a joint survey team from Istanbul University and the University of Chicago crossed the region around Şanlıurfa. The American archaeologist Peter Benedict, already an experienced field researcher in Turkey, recorded large quantities of flint at Göbekli Tepe and recognized a prehistoric site. Some of the stones projecting from the ground, however, were interpreted in relation to later burials. The hill was documented and then left without systematic excavation for three decades (Dietrich, 2016).

That changed in 1994, when Klaus Schmidt arrived. Schmidt, born in 1953 in Feuchtwangen in Franconia, had studied prehistory and protohistory, classical archaeology, geology and palaeontology and had already spent years on excavations in Turkey. Particularly formative was his work at Nevalı Çori, an early Neolithic site on the Euphrates where unusual T-shaped pillars and large-scale sculptures had been found. Nevalı Çori had a bitter distinction: the site was investigated in connection with the construction of the Atatürk Dam and subsequently disappeared permanently beneath the reservoir. The experience mattered when Schmidt visited Göbekli Tepe in 1994, because he recognized in the visible stones forms familiar from the site that had been lost to the reservoir. The 1963 survey team had not failed to see the place. It had simply not yet known what it was looking at. Thirty years later, the same stones were still there, but Schmidt brought with him an experience that had not existed before.

Systematic excavations began in 1995, and beneath the surface appeared large circular and oval enclosures with monumental T-shaped limestone pillars, some around 5.5 metres high and decorated with reliefs of foxes, snakes, wild boar, birds, aurochs and other animals. Some pillars depict arms, hands, belts and loincloths, which is why the abstract T-shapes are interpreted at least in part as stylized human or human-like beings. The oldest of these monumental structures date to the tenth millennium BCE, with the site conventionally dated from around 9600 BCE onward (UNESCO World Heritage Committee, 2018).

Schmidt devoted almost twenty years of his life to Göbekli Tepe and developed an interpretation that would make the site famous around the world. For him, it was primarily a ritual centre created by hunter-gatherers, perhaps a supra-regional meeting place whose monumental architecture may itself have helped draw larger groups into more permanent forms of association. The thesis elegantly inverted a familiar developmental story: perhaps people did not first invent agriculture, settle permanently and then build temples; perhaps shared ritual and monumental construction were themselves among the social changes that encouraged sedentism. The interpretation was compelling, but subsequent excavation made it considerably more complicated, and Schmidt witnessed only part of that reinterpretation. On 20 July 2014, at the age of just 60, he died unexpectedly while swimming during a holiday on the Baltic Sea. Colleagues remembered a man whose life had for years been closely bound to the hill in south-eastern Anatolia; at the time of his death, large parts of the site remained unexplored.

Afterwards, the picture increasingly shifted. More recent excavations found clearer traces of everyday and domestic activity, making it harder to treat Göbekli Tepe as an isolated sanctuary. The German Archaeological Institute now explicitly describes the earlier “mountain sanctuary” interpretation as part of the site’s research history and places greater emphasis on evidence of more permanent occupation. What had once looked like a temple without a village gradually became a much harder-to-separate combination of dwelling, working, eating and ritual, and objects that seem almost banal beside the monumental pillars—particularly grinding stones—have played an important role in that reinterpretation. The prehistorian Laura Dietrich made these tools the subject of intensive study, combining conventional archaeology with experimental reconstruction and use-wear analysis. Visitors understandably photograph the five-metre pillars. Dietrich was interested in what people had spent hours doing with heavy stones on the ground. The answer appears to have been: processing a great deal of grain.

Dietrich and her colleagues identified thousands of grinding and processing implements as well as evidence of intensive use of wild cereals. Animal remains simultaneously show substantial hunting, particularly of gazelle, whose seasonal concentration may point to larger gatherings. Because correspondingly large permanent storage facilities are absent, the researchers proposed work feasts as one explanation: limited periods of collective labour during which substantial amounts

of food were procured, ground, prepared and probably consumed together (Dietrich et al., 2019).

The resulting picture again becomes less monumental and more social. Behind the famous pillars stand not only unknown artists or ritual specialists but people who grind grain, butcher animals, make tools, quarry stone, move loads and feed other people. Monumental construction therefore cannot be explained only through religion or architecture; it also depended on the practical coordination of labour, food and material resources that made collective work possible.

We do not know who coordinated those activities. Nor can we say with confidence whether Göbekli Tepe had pronounced hierarchies, whether particular people occupied special positions as ritual specialists, or how labour and food were distributed among groups. Marc Verhoeven, a Dutch archaeologist who has worked for decades on ritual, symbolism and the Neolithization of the Near East, was still attempting in 2025 to infer relational structures and possible hierarchies from the positioning of human and animal representations (Verhoeven, 2025). The fact that fundamental interpretations continue to shift almost three decades after systematic excavation began says less about inadequate research than about the difficulty of reading social relationships from stone.

The evidence therefore makes it tempting, but methodologically indefensible, to declare some unknown ritual specialist the first Platform Owner, because we know neither a formal organization nor its decision rules. We have no contracts, minutes or names. Yet we can see that activities converged at this place that a single household could hardly have performed alone. People hunted, gathered and processed food; others worked stone or built structures; knowledge, labour and perhaps ritual roles had to come together at the right time.

The infrastructure visible at Göbekli Tepe was therefore not made only of limestone; part of it consisted of relationships. Shared places, rituals, belonging, reputation and common ideas can coordinate human behaviour even where nobody has incorporated an organization, signed a contract or defined a technical interface. Harari's broad claim that *Homo sapiens* can coordinate unusually large groups through shared stories, norms and institutions is not archaeological proof of how Göbekli Tepe functioned. It does, however, remind us that infrastructure is not always made of things that can be touched (Harari, 2015).

Ohalo II and Göbekli Tepe are separated by many millennia and should not be assembled into stages of a neat progress narrative. Together, however, they reveal something that will recur throughout this research essay: humans alter environments, create new opportunities and dependencies, combine their resources with those of others, and develop rules, expectations or common understandings through which interaction becomes more stable and repeatable.

There is still no reason to infer a Platform Owner, market price or Platform Strategy, but structured interdependence is already visible in the way activities, resources and people had to be coordinated.

Across the following millennia, such systems become spatially larger, institutionally more stable and, for historians, more visible. Repeatedly used paths become trade routes, exchange locations become markets, and social obligations increasingly become formalized rules and contracts. Writing, postal systems and later newspapers ensure that human relationships leave not only traces in the ground but documents.

Our proximity to the people themselves changes as well. At Ohalo II, microscopic polish on a flint blade has to tell us what an unknown person may have done 23,000 years ago. At Göbekli Tepe, we can see the pillars but do not know the name of a single person who raised them.

A few millennia later, advertisements, invoices, newspapers and letters bring individual names much closer to view; one of those names leads to London in 1691, a city in which port, exchange, post, press and coffee houses were bringing together ever larger flows of people, information and risk.

The man encountered there, Edward Lloyd, was neither a shipowner nor a captain and, at least initially, probably not even an insurer; he sold coffee.

CHAPTER 3

Edward Lloyd

LONDON · 1688

The coffee was the product. The position was the opportunity.



ON the 10th instant a middle sized Man, having black curled Hair, Pockholes in his Face, an old brown riding Coat, and a black Beaver Hat, was suspected to have taken away 5 Watches, one a pinn'd Case, Silver Box, and a Silver Dial Plate, the Hours cut upon Harts; and a Five Wheel Chain, made by Wilkins of Leicester. The 2d a plain Silver Box with a Glass, the Dial had a pot of Flowers, made by William Corder in Darby. The 3d a Silver Box with a close Silver Case, a pierced Dial Plate with the Day of the Month. The 4th had a Silver Box and a pinn'd Case, many of the Pins being come out, so that the Brails was seen. The 5th had a Silver Box and pinn'd Case, long Hours of the Dial Plate, and Frosted, with a 5 wheel Chain. Whoever gives Notice of them to Mr. Edw. Loyds Coffee House in Tower Street, or to Mr. Edw. Bransby in Darby, shall have a Guinea Reward.

Figure 3.1. The room and the documentary trace: Lloyd's Coffee House as a meeting place for information, commerce and maritime exchange.

RECOGNIZE THE PLATFORM AROUND YOU

The earliest surviving written reference to Edward Lloyd's Coffee House is not about ships or insurance but about five stolen watches.

In February 1688/89, the *London Gazette* carried an advertisement describing a middle-sized man with black curled hair, pockmarks on his face, an old brown riding coat and a black beaver hat, suspected of stealing five watches from an Edward Bransby in Derby. Anyone who could provide information was promised a guinea reward. The notice directed them either to Bransby himself or, in the wording of the advertisement, to "*Mr. Edward Loyd's Coffee House in Tower Street*" (The London Gazette, 1689; Wright & Fayle, 1928).

It is a remarkably unspectacular first trace of a name that would, centuries later, be known around the world for insurance. Lloyd's Coffee House appears neither as a financial institution nor as a marine exchange, but simply as a place where information could be expected to find its intended recipient. Before Lloyd's became a market for risk, it was already something else: an address for information.

We know surprisingly little about the man behind that address. Edward Lloyd was probably born around 1648, and even his place of origin is not entirely certain. Modern biographical work and Lloyd's own historical reconstruction associate him with Canterbury, while older accounts have occasionally described him as Welsh. His biography becomes firmer when he appears in London around 1680 with his wife, Abigail. Nothing yet suggested the later founder figure attached to one of the world's best-known markets (Palmer, 2007; Lloyd's, 2013).

London, however, offered unusually favourable conditions for someone who understood how valuable a meeting place could become. By the 1660s there were already dozens of coffee houses inside the old city walls, and they were far more than places serving a new and bitter drink. Merchants, writers, politicians, sailors and speculators gathered there, read news, received post, talked and did business. Some houses developed recognizable specializations. Anyone wanting to know what was happening in a particular trade or political circle increasingly knew which coffee house to visit (Lloyd's, 2013; Royal Museums Greenwich, 2023).

The English Crown had already noticed that such spaces could matter politically. At the end of 1675, Charles II attempted to suppress coffee houses because, in the language of the proclamation, they were associated with "*false, malicious and scandalous reports*" concerning government and public affairs (Charles II, 1675). The order, which also covered tea, chocolate and sherbet, was intended to take effect in January 1676 but was largely withdrawn before implementation after strong resistance (The Gazette, 2016). A decade before Lloyd's name first entered the surviving record, the government had already recognized that coffee houses did more than sell beverages. They concentrated people and information in ways that were difficult to control.

Lloyd opened his establishment in Tower Street, a location that was hardly accidental. Between the Tower of London, Thames Street, the Custom House, the Navy Office and the river landings moved sailors, captains, merchants and people whose incomes depended directly on what happened at sea. In this part of the City,

ships and money were unusually close. Anyone who needed to meet a captain, sell a cargo or discover whether a vessel had reached a particular port depended on information that was neither fast nor reliably available from a central source (Wright & Fayle, 1928; Lloyd's, 2013).

Lloyd sold coffee, but a good coffee-house proprietor also sold a position within this flow of information.

In 1691 he made a move suggesting that the business was prospering. He left Tower Street and took premises in Lombard Street, only a short walk from the Royal Exchange and the General Post Office. The building was later numbered 16. Wright and Fayle identified it in 1928 through old property records and found an appealing detail: before Lloyd, the premises had for a time housed Pontack's Head, an unusually fashionable French dining establishment for London, where meals could cost several shillings and elaborate dinners one or two guineas. A location associated with expensive food thus became a coffee house whose growing value lay less in the drink than in the people brought together inside it (Wright & Fayle, 1928).

The location was excellent. So was the competition. Around the Royal Exchange stood Garraway's, Jonathan's, the Jamaica, the Marine Coffee House and other establishments with established clienteles. Lloyd did not enter an empty market. He needed to give people a reason to return regularly to his room rather than somebody else's.

Part of the answer was mundane and therefore important: the place worked. Historical reconstructions describe a large, comparatively bare coffee room with a sanded floor and the tables, benches and practical writing materials needed for conversation and business. Lloyd's establishment also employed staff to serve coffee, tea and other drinks. The physical details matter because the later market did not begin as an abstraction; it began as a room in which people could repeatedly find one another, exchange information and transact (Wright & Fayle, 1928; Lloyd's, 2013).

In a world without smartphones, Bloomberg terminals or even a dependable daily information service, a table equipped with paper and ink, surrounded by other merchants and newly arrived news, was infrastructure.

Lloyd began to improve that infrastructure systematically. He gathered news about ships, cargoes and events abroad, used correspondents and made his coffee house a place where maritime information arrived regularly. The Hudson's Bay Company provides a particularly tangible piece of evidence. In 1693, its committee paid "*Mr. Loyd the Coffeeman*" three pounds for intelligence concerning the company's ships. Company records also place the sale of the vessel *Supply* through Lloyd's Coffee House in the same period, for £2,200 (Rich, 1958).

Three pounds for information seems almost trivial beside a ship selling for £2,200, yet for this argument the smaller number is arguably the more revealing one because somebody was willing to pay Lloyd explicitly for knowledge.

The coffee house also became an auction room. Wine, coffee, ships and cargoes could be sold there; around the turn of the century, sales of ships "by candle" are documented at Lloyd's. The uncertain ending of a candle auction made waiting until

the final instant strategically risky. Historical accounts record a wide range of lots passing through the room, from ships to wine and other traded goods (Wright & Fayle, 1928; Lloyd's, 2013).

Lloyd should therefore not be understood merely as a coffee-house proprietor who happened to attract merchants. He increasingly developed services that tied his clientele to the place: news, correspondence, auctions, working facilities and the physical presence of other people with whom business could be done.

In 1696 he went one step further and began publishing Lloyd's News. The paper appeared several times a week and carried domestic, foreign and maritime information. A surviving issue dated 15 October 1696 already shows how far Lloyd had moved beyond selling a beverage: his own name stood above an information product that could circulate outside the room itself (Lloyd, 1696).

The publishing experiment did not last long. In February 1697, *Lloyd's News* reported that the Quakers had petitioned the House of Lords to be excused from all public offices. The Lords declared the report false or misleading and demanded a correction. Contemporary accounts indicate that Lloyd did not simply print the correction in the requested form. Instead, readers were informed that "*Mr Lloyd will print no more at present.*" He stopped publishing the paper (Lee, 1893; Lloyd's, 2013).

Whether this was a matter of principle, commercial prudence or simply the realization that publishing a newspaper could produce far more trouble than serving coffee is impossible to know. The episode nevertheless gives us one rare personal glimpse of an otherwise fragmentary figure. The coffeeman had pushed his information business far enough into public communication to come into conflict with the House of Lords.

The idea did not disappear with the newspaper. Information remained central to the business, and decades later, under different owners, it developed into Lloyd's List, one of the world's longest-running maritime news publications.

The degree to which Lloyd's Coffee House had itself become part of London's public life is captured particularly well by Joseph Addison. In 1711, Addison described in *The Spectator* how he accidentally dropped a sheet containing his private notes at Lloyd's Coffee House. Other customers found it, read it with growing amusement, and finally sent one of the house boys up to the pulpit used for auctions so that he could read the notes aloud to the entire room. Addison, sitting there anonymously, saw only one escape. When the paper eventually reached him, he showed no reaction, rolled it up and used it to light his pipe (Addison, 1711).

The episode may be the most vivid surviving snapshot of the coffee house during Lloyd's lifetime. It is not a solemn insurance market. It is a noisy room full of people, news, paper, auctions, drinks, mockery and business. The same pulpit from which goods were auctioned could, within minutes, become an improvised stage for public information.

At this point, however, an historical correction becomes essential. The modern institutional history of Lloyd's understandably likes to draw a line from the coffee

house to the modern insurance market. The sources from Edward Lloyd's own lifetime support such a linear story only in part.

Marine insurance existed in London long before him. Wealthy merchants accepted portions of individual risks, brokers connected insured parties with capital providers, and coffee houses were naturally suitable places for such deals. It is therefore entirely plausible that insurance contracts were arranged in Lloyd's rooms. What is missing is direct evidence that Edward Lloyd himself was an underwriter, or that during his lifetime the coffee house already functioned primarily as an organized insurance market. A 1928 *Nature* review of Wright and Fayle's major history emphasized that Lloyd himself had little or no knowledge of underwriting and no direct connection with the business; rather, his coffee house became the place increasingly frequented by maritime merchants and insurers (Wright & Fayle, 1928; *Nature*, 1928). More recent legal histories go so far as to note that no unambiguous surviving evidence of underwriting at Lloyd's exists before the mid-eighteenth century, even though such business may well have taken place there during Lloyd's lifetime (Hodgson, 2014).

That evidence changes the story fundamentally because Edward Lloyd did not become important by inventing marine insurance; marine insurance predated him. His historical significance instead lies in the way his room became useful to people who owned ships, commanded ships, traded goods, sought information, assessed risk or needed to transact with one another. The value lay in the density of relevant actors and relevant information.

The more maritime people gathered there, the more useful the house became to other maritime people. The better the information, the greater the reason to come; the more relevant people came, the more valuable the information, auctions and commercial opportunities became in turn. Lloyd did not have to call this a "network effect," or possess a platform theory, to benefit from the mechanism.

His own social position changed as well. Lloyd became a sidesman at St Mary Woolnoth and later held local civic responsibilities. The coffee man who had arrived in London around 1680 had become a prosperous and locally established citizen of the City (Palmer, 2007; Lloyd's, 2013).

His final weeks moved quickly. Early in 1713, Lloyd arranged his will and transferred the coffee-house lease to his head waiter, William Newton, who married Lloyd's daughter Handy shortly afterwards. Edward Lloyd died on 15 February 1713 and was buried at St Mary Woolnoth (Palmer, 2007; Lloyd's, 2013).

Even his death produced one last small trace in the information system of his time. Newspaper notices of private citizens' deaths were far from routine, yet the *Flying Post* reported simply:

"On Sunday last Died Mr Lloyd the Coffee-Man in Lombard Street." (*Flying Post*, 1713, cited in Lloyd's, 2013).

Apparently nothing more was required. The name and the place had already become explanation enough.

Lloyd could hardly have foreseen what followed. The institutional organization of the insurance market emerged only decades after his death; the name survived

successive proprietors, a split into “Old” and “New Lloyd’s,” moves to other premises and eventually the development of a formally regulated market. Particularly revealing is that merchants and underwriters continued to use the name Lloyd’s long after no member of his family remained involved. What had begun as the name of a coffeeman became the name of a place, then an information source, and eventually a market.

The success story also has a side that should not disappear inside a romantic narrative of coffee and entrepreneurship. British maritime trade in the eighteenth century was deeply entangled with colonialism and the transatlantic slave trade, and the later Lloyd’s market became an important insurer within those commercial structures. Lloyd’s publicly acknowledged and apologised for that history in 2020. The institutional peak of this role occurred largely after Edward Lloyd’s death and should not be attributed to him personally; it would be equally misleading, however, to detach the later system from the economic structures whose risks it insured (Oldham, 2007; Lloyd’s, 2020).

That is one reason a sober account of Lloyd is more interesting than a founder legend.

He owned no ships. He carried no cargo. He did not invent marine insurance and, on the available evidence, was probably not himself a major underwriter. What he controlled initially was something far less conspicuous: a room.

Yet into that room came captains, merchants, shipowners, news, letters, auctions and eventually people willing to assume risk. The individual elements already existed elsewhere. Lloyd did not necessarily create them; he brought them into productive proximity.

The strategic point of the case is therefore not that coffee somehow became insurance, but that the coffee house gave Lloyd a position from which increasingly valuable interactions could be organized. He did not need to own the maritime system to occupy an extraordinarily valuable position within it. He first had to recognize who was sitting at his tables, what information they needed, and why returning to his room was becoming more useful for them over time.

Almost two centuries later, another organization would discover that an existing infrastructure could do more than the transaction for which it had originally been built. This time, however, the information would travel not through London by captain, letter or coffee-house waiter, but through a wire.

CHAPTER 4

Western Union

UNITED STATES · NINETEENTH CENTURY

The wire carried no money. The network could still enable the transaction.




Figure 4.1. Person and infrastructure: Hiram Sibley and the operating network that made transactions across distance possible.

RECOGNIZE ANOTHER TRANSACTION YOUR NETWORK CAN ENABLE

In a photograph from 1875, dozens of people sit at long rows of workstations inside the General Operating Department of the Western Union Telegraph Company in New York. In front of them are instruments, paper and the technical equipment through which telegrams from across the country arrived and were sent onward. On one wall is the complex field of connections that a second surviving photograph labels simply “Switch.” Both images are now held by the Library of Congress; the larger photograph carries the note that there are no known restrictions on publication (Library of Congress, 1875a,b).

The photograph is silent, but the room itself formed part of a machine. A telegram submitted in a distant city had to move through wires, relay stations, switches, operators and local delivery staff before reaching its destination. Electrical transmission was only one part of the service. Just as important was an organization that knew which line led where, which operator had to forward which message, and how thousands of individual signals could become a national communications system.

Twenty-four years earlier, such a photograph could hardly have been taken because that system did not yet exist.

By the early 1850s, more than fifty separate telegraph companies were already operating in the United States. They used competing technologies and built overlapping local and regional lines. Telegraphy had solved the technical problem of rapid long-distance communication while creating a new organizational problem: how could numerous lines become a network that worked reliably across much greater distances? (National Park Service, 2024; Wolff, 2013).

One of the men drawn into this disorder had originally had little to do with communications technology. Hiram Sibley, born in Massachusetts in 1807, had built a career around machinery and business before entering telegraphy. By 1849 he was a machine-shop operator, banker, real-estate agent and Sheriff of Monroe County, New York. The man who would later help build one of the country’s largest communications networks did not begin his career with a wire (National Park Service, 2024; Wolff, 2013).

His first venture in telegraphy was not an immediate success. Together with Judge Samuel Selden, Sibley attempted in 1849 to establish a printing-telegraph company. They found themselves in a field crowded with small competitors and concluded that consolidation, rather than another isolated line, offered a different route forward (National Park Service, 2024).

The New York and Mississippi Valley Printing Telegraph Company became one starting point for that consolidation. Sibley’s company acquired additional lines and, in 1856, adopted a name that almost programmatically expressed the direction of travel: Western Union. “Union” was not decorative branding. It described what the business was doing: combining previously fragmented western telegraph lines into a larger system (National Park Service, 2024; Wolff, 2013).

The consolidation was commercially successful but far from politically neutral. Joshua Wolff treats Western Union as a central case in the emergence of large American

corporations during the nineteenth century. Network effects made competition difficult in many telegraph markets, and the entrepreneurs behind Western Union deliberately pursued acquisitions and mergers. By 1866, historical estimates suggest that the company controlled around 90 per cent of the US telegraph market. It also became an early symbol of a new form of corporate power against which politicians, competitors and sections of the press repeatedly protested (Wolff, 2013).

Before Western Union reached this commanding position, however, Sibley took a risk that many other investors preferred to avoid.

California had become increasingly important economically and politically since the Gold Rush, but communication with the eastern United States still depended on slow physical transport. A transcontinental telegraph promised to collapse that distance almost instantaneously. Building it meant running lines across plains, mountains and deserts, spanning enormous gaps between settlements and accepting high construction and maintenance costs.

Even Sibley's own board was initially reluctant. The National Park Service reports that he proposed a transcontinental line to Western Union, but other directors hesitated because of the enormous cost and financial risk (National Park Service, 2024).

In 1860, the US Congress therefore passed the Pacific Telegraph Act, which President James Buchanan signed on 16 June, and the government offered support for construction of a link between the Atlantic and Pacific. Several parties initially showed interest, but the other bidders withdrew, leaving Hiram Sibley in the unusual position of winning the contract partly because nobody else still wanted it.

That did not make the project less risky. The route crossed plains, mountains and deserts and required enormous quantities of poles, wire, insulators, transport and labour. Construction proceeded westward from Omaha and eastward from California through companies organized on Western Union's behalf. On 24 October 1861, the transcontinental telegraph was completed at Salt Lake City, connecting the American coasts electrically (National Park Service, 2024; Wolff, 2013).

The effect was immediately visible. The Pony Express, whose riders had only eighteen months earlier begun carrying mail across the West in spectacular relays, ceased operations within days of completion of the telegraph link. Information that had taken days on horseback could now travel in minutes.

Sibley had helped build a communications network; the strategically more interesting question was what else that network could do once nationwide connectivity, offices and organizational routines were already in place.

When William Orton took over the leadership of Western Union, his own career had been no straighter a path into telegraphy. Born in 1826 in the small town of Cuba, New York, Orton trained first as a printer, worked as a teacher, bookseller and publisher, and later moved into politics. During the American Civil War he entered the newly created federal tax administration; for a time, as Commissioner of Internal Revenue, he led an agency that had been established only a few years earlier to organize the federal government's sharply increased financing needs. He

then moved to the United States Telegraph Company, became vice-president after its merger with Western Union and, in 1867, president of the combined firm. In parallel he studied law and was admitted to the bar in the same year (Western Union Telegraph Company, 1995; Reid, 1879).

Orton was therefore neither an inventor nor a classic telegraph engineer. His career ran through printing, books, administration, taxation and law before reaching telecommunications. Perhaps that combination suited a company whose most difficult problems had become at least as organizational and political as they were technical.

Under his presidency, Western Union reached a scale that was itself politically contentious in the United States. Orton had to manage not only lines and employees but repeated congressional proposals to place telegraphy under state control in a manner similar to the postal service. In 1870, Western Union even published its arguments against a Postal Telegraph System as a separate document running to more than fifty pages (Orton, 1870). Within that document, a short passage is unusually important to the present argument: Western Union explained that it had *“long done something in the way of transmitting deposits and money orders by telegraph”* (Western Union Telegraph Company, 1868). In urgent cases, money orders were already being coordinated through the telegraph network. A customer could deposit money in one place; the information would be sent by telegraph to another, where a corresponding payment could be made. Western Union did not yet promote this as a mass-market service. The company itself gave a practical reason: a fully established system could be comparatively easy for *“rogues”* to abuse. The activity was therefore largely restricted to small sums and urgent circumstances in which a postal money order would arrive too late (Western Union Telegraph Company, 1868).

The passage prevents the history from collapsing into an invention myth, because Western Union did not suddenly discover in 1871 that a telegraph network could also transfer money. The possibility had already been recognized, tested and deliberately constrained; what changed was that an unusual secondary use of existing infrastructure eventually became important enough to be organized as a regular service.

In 1871 Western Union formally introduced its consumer-to-consumer money-transfer service. The year is confirmed not only by later corporate histories but also by US Securities and Exchange Commission documentation and the Department of Justice (Western Union Company, 2007; U.S. Department of Justice, 2005). Later accounts name New York, Chicago and Boston as the first major connections before expansion continued during the same year, although the contemporary primary-source base for that detail is much thinner than for the 1871 introduction itself.

Technically, the decisive mechanism was almost disappointingly simple because the money itself did not travel through the wire. A customer paid money into a Western Union office. Information about the amount, recipient and payout was transmitted by telegraph to another office, where Western Union could pay the recipient from local funds. The physical movement of the exact banknotes handed

over by the sender was irrelevant to the immediate transaction. What mattered was that the company was present in both places, that the two offices could trust one another and that a message could move between them quickly enough.

From the customer's perspective, money had moved from one location to another, whereas from the network's perspective the immediately transmitted object was primarily trusted information. Today this sounds obvious because electronic payment networks have operated on related principles for decades. In 1871, however, it represented a striking extension of what a telegraph company could be. Western Union already possessed lines, operators, offices, accounting systems and a brand that customers trusted to transmit messages accurately across great distances. To create money transfer, the company did not need to build a second transcontinental infrastructure. It needed to organize a different transaction on top of the same infrastructure.

The surviving corporate records show how extensive the organization behind that apparently simple idea became. The Smithsonian now holds 452 cubic feet of Western Union material in 871 boxes and 23 map folders, including correspondence, operating material, annual reports, financial records and presidential letterbooks. The archive makes the network visible as something much larger than wires: it was also rules, employees, tariffs, contracts and countless organizational decisions (Western Union Telegraph Company, 1995).

Western Union's own early concern about fraud also shows that trust and governance were at least as important to money transfer as electrical speed. A communications network could transmit an instruction rapidly. It also had to ensure that the right person received the right money at the other end. Surviving operating and money-order records document the procedures through which the company tried to turn electrical communication into a controlled financial service (Western Union Telegraph Company, 1995).

Western Union had therefore done more than add a second product to a price list: the same geographic presence, technical connectivity and trust structure that supported a message network could also be organized as transaction infrastructure.

The difference is subtle but strategically important. In a telegram, the message is itself the thing the customer wants transported. In a money transfer, the message becomes the mechanism through which something else is coordinated. The telegraph was no longer only the product. It became infrastructure for a transaction.

William Orton had only a limited time to see the long-term value of that extension. On 22 April 1878, he died in New York at the age of 51. In the surviving corporate archive, his presidential correspondence is followed by that of his successor, Norvin Green. In a company whose business was the uninterrupted forwarding of information, the leadership itself passed to the next president (Reid, 1879; Western Union Telegraph Company, 1995).

Western Union remained known primarily as a telegraph company for many decades after that, but over the longer historical arc the telegraph disappeared from the centre of the business while money transfer survived.

There is perhaps no more striking irony in the case. The technical infrastructure for which Western Union had become famous was largely displaced during the twentieth and early twenty-first centuries by telephone, digital networks and internet communication. The transaction the company placed on top of that infrastructure in 1871, by contrast, eventually became its core business. Modern corporate filings still identify 1871 as the beginning of Western Union's consumer-to-consumer money-transfer service (Western Union Company, 2007).

Sibley had tried to turn many lines into a network. Orton's Western Union later recognized that an existing network could carry more than one kind of transaction.

For Platform Awareness, it matters less whether we can reconstruct which individual employee first proposed in 1871 that occasional telegraphic money orders should become a systematic service. The surviving sources do not yet support that personal invention story, and it would be wrong to attribute to Orton a documented flash of insight for which we have no evidence.

The strategic movement is visible without a hero story because Western Union already possessed a national communications network. The next move was not simply to build more of the same. It was to recognize that the same position between geographically distant people could enable a different transaction.

The wire carried no banknotes, but it made money available somewhere else, illustrating how a new platform position can emerge not from building an entirely new infrastructure but from recognizing what additional transaction existing infrastructure can support.

The next case changes the logic again because the company owns neither the communications network nor the most important physical infrastructure on which its growth depends; its advantage comes instead from learning how to organize a business on top of them.

CHAPTER 5

Richard Warren Sears

UNITED STATES · LATE NINETEENTH CENTURY

Building on other people's infrastructure.




Figure 5.1. Entrepreneur and offer: Sears scaled a mail-order model through postal and railway infrastructures he did not own.

RECOGNIZE THE PLATFORM YOU DO NOT OWN

Sometime in the mid-1880s, a parcel was left sitting at a small railway station in Minnesota. Nobody there seemed to want it.

The station was at North Redwood, near Redwood Falls, in a region where railways did more than move people and goods. They connected remote places to larger markets with a reliability that had not previously existed. One of the men working at this node was Richard Warren Sears, in his early twenties, a telegraph operator who also handled express shipments. Sears had been born in Stewartville, Minnesota, in 1863 and had followed no course of education that pointed toward a future as one of the best-known retail entrepreneurs in the United States. As a young man, he earned his living on the railway and for a time also ran a small business in coal, wood and lumber. Even in 1907, by then a millionaire, he told a *Chicago Tribune* reporter that without one particular watch he might still be somewhere in Minnesota dealing in coal, wood and lumber (Chicago Tribune, 1907).

The story of that watch has been told in several versions over the decades, which is precisely why it should not be made too neat. A near-contemporary account in *Printers' Ink* in 1914 reported that a watch company in the East had sent a shipment cash on delivery to a local jeweller. The jeweller refused to accept it. Sears, who as express agent had to deal with the unwanted parcel, wrote to the sender and offered to sell the watches himself on commission. The company agreed. Sears then contacted railway workers in the surrounding area by letter and telegraph and sold the shipment quickly (Printers' Ink, 1914).

The *Chicago Tribune* account published in 1907 tells the opening somewhat differently. There, the story begins with a single silver watch sent to Sears from Fredonia, New York, which revealed to him how large the gap could be between a product's price at source and the price at which it normally sold in the North-West. The newspaper gives a purchase price of about \$9.40 for a comparable watch and a possible selling price of \$11.90, while railway workers traded similar watches among themselves for around \$20. Sears then supposedly wrote forty letters to men in forty places and made about \$200 on the first attempt (Chicago Tribune, 1907).

It is difficult today to establish with complete confidence which version most accurately captures the first transaction. Even the year is given as 1884 in some later accounts and 1886 in others. The Library of Congress dates the founding of the R. W. Sears Watch Company to 1886 and treats the story of a shipment rejected by a local jeweller as the company's origin (Library of Congress, 2026).

For this argument, however, the precise number of watches matters less than the position from which Sears saw them. He was not in a major city and did not own a national retail network; he was standing at a railway station where several systems converged. The railway brought products from distant manufacturing centres. The telegraph connected station employees to one another. Express companies enabled cash-on-delivery shipments. The postal system allowed Sears to reach people he had never met. As a railway employee, he also knew a group of potential customers distributed along the network, whose addresses and communication routes were comparatively easy for him to reach.

The watch was merchandise, whereas much of the leverage that allowed Sears to find and serve dispersed customers lay in the infrastructure around it.

Sears initially sold watches on the side. Later accounts say he still wrote his advertising letters by hand and enlisted another railway employee only when the correspondence became too extensive (*Printers' Ink*, 1914). The business soon became more profitable than his railway job. He left his post, moved to Minneapolis and turned the side business into his occupation.

The path was by no means a straight line toward the later Sears, Roebuck & Company. Sears sold his first business, withdrew temporarily from the Chicago market and then started again. During this phase Alvah Curtis Roebuck, a watchmaker, joined him and supplied a repair capability to a business whose founder's greatest strength was selling. Sears and Roebuck eventually moved their activities to Chicago, and watches and jewellery gradually expanded into a much broader assortment (Library of Congress, 2026).

Sears himself seems to have had little retrospective interest in turning his biography into a story of prophetic genius. The *Chicago Tribune* reporter who visited him in 1907 found a very wealthy man who still attributed his rise to contingencies and to other people. Sears pointed to the watch that had been sent to him and to an older merchant named Wynkoop, from whom he said he had learned to look at goods from the customer's perspective. The newspaper nevertheless rendered his rise almost as an American Aladdin story: instead of a magic lamp, Sears had received a silver watch (*Chicago Tribune*, 1907).

Nor did he appear particularly interested in performing the role of the conventional captain of industry. At his summer property near Grayslake, Illinois, the same reporter described Sears wearing a flannel shirt and overalls while his four children ran around bareheaded and suntanned. One local anecdote seemed especially characteristic to the journalist: Sears supposedly allowed a child to spend only ten cents out of fifty on some small purchase and then, immediately afterwards, paid fifty dollars for two dogs that one of his children desperately wanted. Whether this demonstrated frugality, contradiction or simply a father with his own ideas about sensible spending was left largely to the reader (*Chicago Tribune*, 1907).

The relative informality described at Grayslake contrasts with the aggressiveness of Sears's commercial practice. His particular talent did not lie in inventing mail order. Aaron Montgomery Ward had begun selling goods by catalogue to rural customers in 1872. Sears entered an existing market and pushed its communication practices with unusual aggression (National Postal Museum, 2016).

He grasped early a problem that distinguishes distance selling from an ordinary store. A customer who cannot touch the product and has never met the seller has to be persuaded in some other way to send money to a faraway place.

Sears tried almost obsessively to reduce that distance through language, and a retrospective in *Printers' Ink* in 1914 described how minutely he explained products in advertisements. Clothing descriptions should cover not only fabric and cut but even lining, buttons and buttonholes. His famous "send-no-money" advertisements

violated, in the recollection of one industry observer, almost every contemporary rule of good advertising: small type, densely packed lines, little visual elegance and unusually long descriptions. Other advertising professionals considered them barely readable. Sears's response was not to make the advertisements prettier, but to make the content interesting enough that people would read the small print anyway (Printers' Ink, 1914).

He also reduced the psychological threshold for ordering. Customers were encouraged to write in their own words, with pencil or pen, on an order form or, if necessary, on any scrap of paper. If they disliked a purchase, Sears promised generous return terms. The catalogue therefore tried to replace the shopkeeper who would normally explain, reassure and create trust.

The catalogue became an interface between a company and people living hundreds or thousands of kilometres away.

For this to work, Sears needed to understand how those people actually responded. An unusually rich source is the Louis E. Asher Papers at the University of Chicago. Asher, originally a reporter and advertising copywriter, joined Sears at the end of the 1890s and later became General Manager. The archive contains approximately 200 letters, including 120 from Richard Sears to Asher, together with advertising material, customer correspondence, cost calculations and memoranda on promotional ideas. They show not an entrepreneur who simply outsourced advertising to an agency, but an organization that made language, response and customer behaviour part of everyday operating work (University of Chicago Library, 2006).

Sears could be almost pedantic about it. A later colleague recalled how he would reject catalogue pages even after they had been set because a product did not look right. With a watch, he supposedly used his hands in the air to demonstrate the angle at which it ought to appear. The same witness captured this side of Sears in an unexpectedly elegant phrase: *"He had the soul of an artist"* (Printers' Ink, 1914).

The meticulous treatment of advertising could not by itself move millions of parcels, and as the business grew, Julius Rosenwald became increasingly important. His personality differed markedly from that of the restless promoter Sears. The son of Jewish immigrants in Springfield, Illinois, Rosenwald had worked in the clothing trade before joining Sears, Roebuck as vice-president in 1895. He later became President and then Chairman. Under his leadership, the organization, reliability and efficiency of the mail-order operation were systematized in particular (University of Chicago Library, 2009; Emmet & Jeuck, 1950).

Sears and Rosenwald complemented one another, but their differences also generated friction. Sears was an aggressive advertiser and growth-oriented entrepreneur; Rosenwald was regarded as financially more cautious. When the financial crisis of 1907 hit a company that had just invested heavily in an enormous new Chicago complex and further expansion, Sears wanted to answer declining orders with still more advertising. Rosenwald considered that too risky. In 1908 Richard Sears stepped down as President. The man whose name appeared on millions of catalogues

therefore disappeared from the operational centre of a company that would continue to grow for decades (University of Chicago, 1949).

By then, Sears had already built an enormous physical machine of its own.

Beginning in 1905, a mail-order complex rose on Chicago's West Side. Its central Merchandise Building was later described by the Historic American Buildings Survey as a vast, carefully coordinated machine for receiving, storing, assembling and dispatching goods. It was not merely a warehouse. Orders, information and merchandise moved through the complex via a dense system of routes, departments, conveyors and pneumatic tubes (Historic American Buildings Survey, 1933).

Library of Congress materials still preserve plans of the conveyor systems and photographs of the pneumatic-tube stations. In one surviving image, a row of tubes stands against a wall like a technological nervous system. Other photographs show conveyors running between storage racks or dispatch areas where streams of merchandise were brought together. The HABS documentation also preserves detailed drawings of the building and its internal systems (Historic American Buildings Survey, 1933).

This internal infrastructure belonged to Sears, whereas two of the infrastructures that made national scale possible did not, beginning with the railway. Chicago had become one of the most important railway hubs in the United States. Sears did not need to build its own track to Minnesota, Iowa, Nebraska or Texas. It needed to organize fulfilment so that goods could enter existing rail networks. The very infrastructure on which Richard Sears had encountered his first opportunity as a young telegraph operator therefore remained essential to the later scaling of the company.

The other was the US postal system, where public infrastructure dramatically expanded the addressable market. With Rural Free Delivery, introduced from 1896, mail increasingly reached rural households directly instead of forcing residents to travel long distances to a post office. Catalogues, orders and payments could therefore circulate far more easily between Chicago and farms or small towns. Mail-order commerce and postal expansion reinforced one another: more catalogues generated orders; orders generated parcels; improved delivery in turn made the catalogue more attractive (National Postal Museum, 2016).

Sears did not have to invent Rural Free Delivery or persuade the US government to build a dedicated "Sears network"; the strategic advantage came from fitting its business model increasingly well to infrastructure that already existed and whose governance remained outside the firm.

In 1910, according to the Smithsonian, Sears, Roebuck mailed 2.3 million catalogues. A rural household could now access hundreds of pages of merchandise that no local general store could ever have held in stock (National Museum of American History, 2015).

A second major infrastructure shift followed in 1913. Until then, postal carriage of heavier packages was sharply constrained, leaving private express companies dominant in much of the parcel business. After years of political argument, the US

Parcel Post Service began on 1 January 1913. At first, packages of up to eleven pounds could be sent through the postal system. Demand was extraordinary: around 300 million parcels were shipped in the first six months (Pope, 2013).

For Sears, the effect on mail order was immediate and substantial. Sears handled five times as many orders in 1913 as in the previous year; within five years, the company's revenues had doubled (Pope, 2013; National Postal Museum, 2025).

That number should not be read as evidence that Parcel Post alone made Sears large. The company was already an enormous mail-order retailer with its own logistics, catalogue production, advertising system and procurement network. Parcel Post instead changed the cost and reach of the connection between that existing system and households beyond Chicago.

The strategically interesting part of the case lies in the relationship between these systems. Inside the firm, Sears increasingly controlled a sophisticated platform of catalogue, data, procurement and fulfilment. Customers on one side could choose from a vast range; manufacturers on the other received access to nationally aggregated demand. Sears organized selection, prices, communication and trust.

Outside this system, however, Sears did not own some of the infrastructure most important to its scale: customer letters travelled through the post, merchandise leaving Chicago moved through railway networks, and delivery into rural America increasingly depended on public postal infrastructure.

Sears owned the catalogue, the assortment, the brand and the internal fulfilment machine. It did not own the transportation and communication environment within which those assets could scale nationally.

Its advantage depended on positioning the business within that environment, a relationship that became so natural that Richard Sears may have understood it less as theory than as practice. His first customer group consisted of railway workers; his first business emerged at a railway station; his advertising travelled by post; his catalogue became more valuable through Rural Free Delivery; his goods could reach deeper into American territory through rail and, later, Parcel Post.

From this perspective, Sears is not only a classic example of a company that built a successful business model. It is also a case of growth through reading somebody else's system better than others did.

Sears did not have to own either a national railway or a national postal system; it had to understand what could be organized on top of infrastructures whose scale and governance remained largely outside the firm.

The most appealing contrast returns us to Richard Sears himself. Millions of catalogues were circulating through a vast information and logistics network, while a mechanical complex on Chicago's West Side moved orders through pneumatic tubes and merchandise along conveyors. Yet a reporter visiting Sears in 1907 described him on his Grayslake farm in a flannel shirt and overalls.

A man could apparently help build one of the most complex commercial machines of his age and still look at home as though he might need to go and check the corn.

Richard Sears died on 28 September 1914 at the age of fifty. *Printers' Ink* soon published a substantial obituary for a man whose advertising practices the journal had observed for years. The company bearing his name continued and, under Rosenwald and later leaders, developed into the largest retailer in the United States (Printers' Ink, 1914).

Yet the original strategic advantage is still visible at that small station in Minnesota.

At that station Sears owned almost none of the wider system that would later make national scale possible—not the railway, telegraph, express network or postal system—but he occupied a position at their intersection from which a commercial opportunity became visible.

With Lloyd, value emerged because relevant actors and information converged in a room. Western Union discovered that a network it owned could support another transaction. Sears reveals a third movement: an organization does not have to own a platform to benefit strategically from its structure.

The case therefore suggests that strategically decisive infrastructure may already exist, in which case the relevant question concerns the position an organization can occupy on top of it rather than whether it should reproduce the infrastructure itself.

The next case moves more than a century forward and reverses the perspective again, because the overlooked platform is not outside the organization but already inside it.

CHAPTER 6

Amazon Web Services

AMAZON · 2000

The platform was already inside.

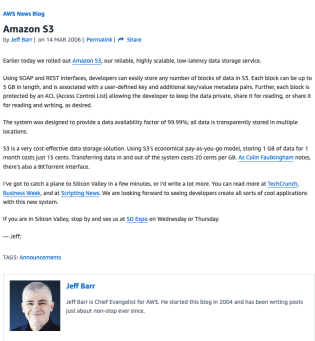


Figure 6.1. Person and document: Amazon CTO Werner Vogels represents the architectural and technical leadership surrounding Amazon’s shift toward service-based infrastructure; Jeff Barr’s AWS News Blog post documents the public launch of Amazon S3 on 14 March 2006. Together they mark the movement from internally developed technical capabilities toward infrastructure accessible to outside developers.

RECOGNIZE THE PLATFORM INSIDE YOUR ORGANIZATION

In the early 2000s, Amazon had an internal process known by the acronym NPI. When a team wanted to build something new and needed support from several other parts of the company, the initiative went through this process. Each participating team would estimate how many people-weeks of work it would need to contribute.

On paper, that sounds reasonable. Inside a rapidly growing technology company, however, it became precisely the kind of coordination mechanism that eventually produces more coordination than innovation.

More than twenty years later, Andy Jassy remembered these meetings in language that betrayed little nostalgia. He described NPI as a kind of “*cabal*”, a heavyweight forum through which new initiatives had to negotiate the interdependencies among teams. The bottleneck constrained what Amazon could accomplish and, as Jassy later put it to shareholders, “*frustrated the heck out of us*” (Jassy, 2024). The important response was not to organize better meetings. Amazon began to dismantle the technical dependencies themselves.

Nothing about Jassy’s arrival at Amazon suggested that he would later build one of the most important infrastructure platforms in the computing industry. He had originally wanted to become a sports broadcaster, sent audition tapes to small television markets, then worked in television production, coached the football team at his old high school and for a time sold golf equipment. Product management followed, then entrepreneurial experiments, brief periods in sales and investment banking, and finally Harvard Business School. In May 1997 he took his last exam there; three days later he started at Amazon, by his own later account without knowing exactly what his job or title would be (Jassy, 2025).

Amazon itself was still far from the company later associated with the name. In 1996 it had generated around \$15 million in revenue, sold books, had no third-party Marketplace and shipped only within the United States. Jassy joined in 1997 and, over the following years, gained a broad view of a company that was rapidly adding businesses and accumulating increasingly complex technical dependencies (Jassy, 2022).

One of those problems emerged from a business now almost forgotten in the shadow of AWS: Merchant.com.

Amazon had begun using its e-commerce technology for large external retailers. Companies such as Target were supposed to use Amazon’s technical components as the backbone of their own websites while placing their own brand and customer experience on top. Strategically, the idea sounded straightforward. Technically, it was not.

Amazon had grown at extraordinary speed during its first years. Capabilities such as search, payments, product information and ordering already existed, but they were tightly entangled. As long as Amazon itself remained the only user, many of these connections could somehow be managed internally. As soon as Target needed to use individual functions, Amazon discovered that its apparent building blocks could not be separated cleanly.

Jassy later described the problem with unusual directness: the components were *“too jumbled together and not partitioned right”* (Jassy, 2024). What looked like a ready-made technological toolkit was in fact a densely coupled system of mutual dependencies.

Target thereby became an unintended stress test for Amazon’s own architecture, because if an external organization was to use Amazon’s search, ordering or payments capabilities, those capabilities needed clear interfaces. A team could not call several other teams every time, learn their internal implementation or negotiate a bespoke integration process. Functions had to be separated, documented and made available through stable application programming interfaces, or APIs.

What Amazon needed for external retailers soon proved useful inside the company as well.

Jassy later recalled that from around 2000 Amazon increasingly expected internal teams to build services in ways that other teams could access through clearly defined APIs. Almost without public attention, he said, Amazon became internally a services company (Miller, 2016).

The service-oriented reorganization solved only part of the problem. Amazon hired more and more software developers and still found that new applications were not appearing as quickly as expected. In Jassy’s later account, project teams kept running into the same reality. An initiative that management imagined as a three-month development effort might spend those same three months merely getting a database, computing capacity or storage in place. Highly skilled developers were devoting substantial time to rebuilding technical foundations that other teams elsewhere inside Amazon needed as well (Miller, 2016).

At first there was no new business proposition, only an internal productivity problem, and that distinction matters. The people who would later build AWS did not begin by asking which entirely new industry Amazon might enter. They were dealing with capabilities that already existed inside the company but were too difficult to access, too tightly coupled or too slow for their internal users to obtain.

Amazon already stored very large amounts of data, handled enormous computing loads and operated highly scaled technical infrastructure; the unresolved question was whether those capabilities could be organized so that developers no longer had to reconstruct them repeatedly.

Around 2002, another signal appeared that the logic might work beyond Amazon itself. The company opened early parts of its e-commerce system to outside developers through web services. One of the people electrified by this was Jeff Barr.

Barr had worked at Microsoft around Visual Basic and had become unusually immersed in emerging technologies such as XML, RSS and web services in his spare time. His side project Syndic8, a platform for collecting and organizing RSS feeds, had grown by his later recollection to almost 90,000 lines of PHP and more than fifty database tables. What looked like a hobby had quietly given him practical experience with the very problems of hosting, scale and databases that would soon sit at the centre of AWS (Barr, 2019).

When Barr discovered Amazon's first web service in 2002, he tried it, sent feedback and, along with several other early users, was invited to spend a day at Amazon. There he heard an employee explain, in substance, that Amazon intended to look across the company for more functionality that could be opened through web-service interfaces. Barr later called it his "light-bulb moment." During the event he told an Amazon employee that he wanted to work on exactly that. A few weeks later, he was working at Amazon himself (Barr, 2019).

This was not yet AWS in the form we now recognize. The early services mainly exposed Amazon product data and e-commerce functionality to outside developers. But they demonstrated something important: people outside the company could take Amazon capabilities and use them to build applications Amazon itself had never imagined.

Experience with the early web services therefore widened the question from which e-commerce functions could be exposed to whether the technical foundations themselves could be offered through similarly accessible interfaces.

By 2003, these ideas were converging. Jassy was working closely with Bezos and began developing a concept for a far more fundamental infrastructure business. A passage from the 2003 AWS Vision document, reproduced by Jassy years later, shows how far the thinking had already developed. Amazon wanted to build primitives: small, fundamental technical building blocks, each doing one thing exceptionally well, as indivisible as practicable and free for developers to combine in ways of their own choosing (Jassy, 2024).

The principle ran almost opposite to the dominant logic of large enterprise-software suites, which bundled as many capabilities as possible into comprehensive solutions. Amazon instead wanted to offer storage as storage, compute as compute and further capabilities as independent building blocks whose eventual combinations would not have to be predetermined by the provider.

In that architecture, the provider did not need to know all the applications that would eventually be built on top of the platform, provided that the underlying building blocks were sufficiently useful, stable and composable for others to invent with them.

How exactly the ideas of 2003 converged is not remembered identically even by participants. Jeff Barr wrote in 2019 about a day-long brainstorming session at Jeff Bezos's house, attended by Bezos, Jassy, engineer Al Vermeulen, Barr himself and several others. The group assembled a long list of possible services from which S3, EC2 and RDS would later emerge. Barr added a healthy caveat to his own recollection: he was fairly sure this was the meeting later described in AWS origin stories, but not completely certain. He even noted that he should have taken better notes (Barr, 2019).

The uncertainty is useful precisely because it resists a tidy founding myth: AWS did not emerge from one isolated eureka moment. It grew out of years of technical frustration, internal reorganization, Merchant.com, APIs, developer feedback, infrastructure problems, written proposals and discussion. Jassy himself later stressed that there was no single magical instant in which cloud computing suddenly appeared

fully formed (Miller, 2016). As the concept became more concrete, the proposed primitives had to become operational technologies, a transition that can be followed unusually closely in the development of the first major infrastructure service, Amazon Simple Storage Service, or S3. The Dutch computer scientist Werner Vogels joined Amazon from Cornell University in 2004 after years of research on distributed systems. He had completed his doctorate at Vrije Universiteit Amsterdam in 2003, and his academic work focused on communication technologies, scalability and reliable distributed computing. That background aligned closely with the problems Amazon was trying to solve as it turned internal infrastructure capabilities into services for outside developers (Cornell University, 2004; Vogels, 2022).

Some early architectural work took place in settings that hardly resemble the birthplace of global cloud infrastructure. The S3 team discussed design principles over several days with Hammerhead Ale at Six Arms, a pub in Seattle's Capitol Hill neighbourhood, and over bagels at the Seattle Convention Center. Hardware failure was to be treated not as a rare exception but as a normal condition in a large distributed system; single points of failure had to be eliminated and components designed so that later services could build on them (Amazon Staff, 2021).

The informality of those settings should not obscure the architectural principle being developed: foundational building blocks become useful to other actors only when they can be relied upon as stable infrastructure.

On 14 March 2006, S3 became publicly available, and the launch was strikingly understated. Jeff Barr published a short blog post explaining that developers could use REST and SOAP interfaces to store any number of data objects, then up to five gigabytes each. One gigabyte of storage cost 15 US cents per month, with data transfer priced at 20 cents per gigabyte. The service was engineered for 99.99 per cent availability and stored data transparently across multiple locations (Barr, 2006a).

Barr's announcement ended almost casually with the remark that he had to catch a flight to Silicon Valley in a few minutes and therefore could not write more, leaving a launch post of only a few paragraphs for a service that would later contribute to a substantial change in the economics of software infrastructure.

Werner Vogels described the strategic significance more fundamentally on the same day. S3, he wrote, was another step in Amazon's mission to expose the "*atomic-level pieces of the Amazon.com platform*" to outside developers (Vogels, 2006). Scalable, reliable and secure storage already available to Amazon's own builders would now be available to developers outside the company.

That formulation captures the decisive movement of the case: capabilities Amazon had developed and standardized for its own purposes were being turned into accessible infrastructure for external users.

S3 on its own initially seemed odd to some developers. Jassy later recalled that users found the service useful but also asked why Amazon was offering only storage. The answer became clearer a few months later (Jassy, 2024).

On 24 August 2006, Amazon announced the public beta of Elastic Compute Cloud, EC2. Developers could now start and stop virtual machines within minutes and

pay only for the computing time they actually used. What had previously required buying, installing and operating servers could be consumed as an on-demand resource (Amazon Web Services, 2006).

There is a strikingly personal turn behind this service as well. Chris Pinkham, who had been responsible at Amazon for areas including global networking and data-centre operations, had a second child and wanted to return with his family to his native South Africa. Amazon wanted to keep him and allowed him to establish a development centre in Cape Town in 2004. There, Pinkham assembled with a small group of predominantly South African developers the team that built important parts of what became EC2. A global compute platform therefore emerged partly ten thousand kilometres from Seattle because an engineer wanted to live closer to his family (Pinkham, 2011).

The story has a fitting small sequel. When Jeff Barr announced the EC2 beta on the AWS blog in August 2006, he too was not sitting in a sterile launch room in Seattle. He wrote from the “steaming hot beaches of Cabo San Lucas”, where he was ostensibly on holiday. For ten cents per hour, developers could then rent a virtual machine roughly equivalent to a 1.7 GHz Xeon processor with 1.75 gigabytes of RAM and 160 gigabytes of local disk (Barr, 2006b).

The casual launch settings should not obscure how consequential the underlying economic model was. A young software company no longer had to buy server capacity years in advance, equip a data centre and hope that its business would later grow enough to justify the investment; it could rent storage and compute in small quantities and add capacity when it actually needed it. For many customers, this shifted part of infrastructure expenditure from fixed investment toward variable cost, made capacity more available on demand and allowed software to request technical resources through APIs rather than through lengthy procurement and provisioning processes.

Jeff Bezos made clear in his 2006 shareholder letter that Amazon now recognized the significance. AWS was a new business serving a new customer set: software developers. At that point Amazon offered ten web services and counted more than 240,000 registered developers. Particularly revealing was Bezos’s explanation of why Amazon believed it could operate such a business at all. Storage and compute were universal developer needs, while Amazon had developed deep expertise in both through roughly twelve years of scaling Amazon.com (Bezos, 2007).

The capability had not first been discovered through a conventional outside-in market analysis; Amazon had developed it for itself and only gradually learned to interpret it as something that could serve an external market. The shift becomes especially clear in the company’s 2007 annual report, which described Amazon around three primary customer sets—consumers, sellers and developers—and stated that AWS provided developers with direct access to Amazon’s technology platform so that they could build applications on top of it (Amazon.com, Inc., 2008). Within only a few years, the same company therefore occupied several simultaneous positions: consumers encountered a retailer and marketplace, sellers encountered a sales

channel and increasingly fulfilment infrastructure, while developers encountered an infrastructure provider. These roles did not replace one another because they belonged to different platform relationships.

The longer-term impact of AWS reached beyond Amazon's own growth. Michael Ewens, Ramana Nanda and Matthew Rhodes-Kropf later studied how the introduction of cloud computing changed the financing of young companies. Coming from finance and entrepreneurship research, they treated AWS not merely as an interesting technology but as a measurable shock to the cost of experimentation (Ewens et al., 2018).

Their analysis of venture-capital financing found that start-ups in industries especially able to benefit from AWS raised on average around 20 per cent less capital in their first funding round after its introduction. Among firms that survived for several years, the difference in total capital raised largely disappeared. The pattern is consistent with an important shift: large investments no longer had to be made in full before anyone knew whether an idea worked. Start-ups could experiment earlier and more cheaply and postpone heavy infrastructure spending until more was known about their chances of success (Ewens et al., 2018).

AWS therefore did more than alter the unit economics of servers; it changed the timing and scale at which customers had to take infrastructure risk.

The Amazon case therefore differs from the preceding cases in the direction from which the strategically relevant system became visible. Edward Lloyd occupied a valuable position in a maritime system he did not own, Western Union recognized another transaction that could be supported by a communications network it did own, and Sears learned to scale through railway and postal infrastructures governed by others. At Amazon, by contrast, much of the platform was already inside the organizational boundary, where databases, storage, compute, payments, search, ordering and fulfilment were initially treated as technical prerequisites for other businesses. Their meaning changed as internal dependencies became painful, external merchants demanded cleaner interfaces and developers began using Amazon capabilities in unexpected ways, making it possible to reinterpret what had looked like cost and internal support as reusable products for external Consumers and, collectively, as a platform of technical primitives.

That is why it is too simple to tell AWS as the story of an online retailer that one day decided to rent servers. Amazon did not start with empty data centres in search of an additional use. It had spent years developing capabilities to cope with its own extraordinary growth, separating those capabilities from one another, making them accessible through APIs and eventually recognizing that the problem repeatedly encountered inside Amazon also existed outside it.

The strategic question of the case is therefore whether the platform an organization is looking for may already exist inside it in a form that has not yet been interpreted as a distinct platform opportunity.

Platform Awareness looks not only outward toward markets, ecosystems or infrastructure owned by others. It also requires organizations to examine capabilities

inside their own boundaries that are repeatedly used by multiple actors, whose value increases through reuse and that might move from internal technical support to an independent platform.

Amazon therefore did more than recognize a new external market; it learned to interpret itself differently. The next case takes that movement one analytical layer further by showing that even when a company already owns and governs a successful platform, its position is not permanently settled, because another platform can form around or above it and require the existing orchestrator to participate in a system it does not control.

CHAPTER 7

Shopify

COMMERCE · 2004–2026

Even a platform can become a participant.




Figure 7.1. Founder and origin: a merchant's own technical solution became infrastructure for other merchants, then a platform inside larger commerce systems.

RECOGNIZE THE PLATFORM ABOVE YOUR PLATFORM

By the spring of 2026, an online purchase no longer had to begin in an online store. A buyer could describe an imprecise need to an AI system, let us say, we are looking for a compact suitcase for a weekend trip, durable enough for frequent travel and priced below a particular threshold, and allow the interface to identify products, compare variants, ask for missing information and, where the service and merchant supported it, participate in checkout. The product might still belong to an independent merchant and the transaction might still rely on the merchant's existing commerce infrastructure, yet the surface on which discovery and deliberation occurred could belong to ChatGPT, Microsoft Copilot, Google's AI Mode or Gemini rather than to the merchant itself (Shopify, 2026a; Google, 2026a). For Shopify this development posed a problem that is easy to miss if a successful platform is assumed to occupy one stable strategic position. By the end of 2025, merchants using Shopify had generated \$378 billion in gross merchandise volume, while Shopify Payments processed \$248.1 billion, and the company had spent two decades integrating storefronts, product data, checkout, payments, applications, point of sale and other commerce functions into what it increasingly described as foundational infrastructure (Shopify Inc., 2026). A new interface layer did not make that infrastructure disappear, but it raised a different question: what happens when a platform that orchestrates its own ecosystem becomes a participant in a larger system whose customer interface it does not own?

There is an historical symmetry in Shopify reaching this point because the company itself emerged from a problem that Tobias Lütke had encountered as a participant in somebody else's technical environment. Lütke, born in Koblenz in 1981, had begun programming as a child, left secondary school after the tenth grade and completed a vocational apprenticeship in software development before working with enterprise systems and open-source software. Shopify's later corporate filings continued to note his involvement in the core team around Ruby on Rails and in open-source libraries such as Active Merchant (Shopify Inc., 2026). When he moved to Ottawa in 2002, initially following his girlfriend, later his wife, he continued working remotely for a German software company but increasingly found the work isolating; snowboarding offered a more tangible interest, and in 2004 he and his co-founders started Snowdevil with the comparatively ordinary intention of selling snowboards and related equipment online (Lütke, 2010; Ottawa Business Journal, 2017). The difficulty was not the idea of the shop but the software through which it had to operate. Existing systems appeared to Lütke expensive, cumbersome and oriented toward established businesses moving online rather than small merchants beginning there, an experience he later summarized in the letter accompanying Shopify's public listing by writing that the software landscape had seemed to work against the founders' ambitions at almost every step (Shopify Inc., 2015).

Lütke therefore returned to programming, sometimes working from a Bridgehead Coffee House on Elgin Street in Ottawa, where he later joked that he had secured access to the Wi-Fi after fixing the café's router (Ottawa Business Journal, 2017). The setting invites an easy comparison with Edward Lloyd's coffee house three

centuries earlier, but the more useful continuity lies elsewhere: both stories begin with a mundane place or product whose value changes as increasingly important interactions become organized around it. In Lütke's case, the technical friction extended from storefront software to the payment interface itself. He later recalled sending notarized copies of his passport to American Fork, Utah, in order to obtain API credentials for accepting credit cards, an episode that made concrete how many institutional and technical steps separated the simple intention to sell something from the completed transaction (Shopify, 2024). Snowdevil ultimately worked, its first sales season was profitable, and Lütke had built much of the commerce software himself in Ruby on Rails. Yet the founders gradually became more interested in the reusable system underneath the shop than in the snowboards sold through it. As Lütke later described the transition, they had built the software they wished had existed when they started; merchant frustration became a product for other merchants (Lütke, 2010; Shopify Inc., 2015). Shopify launched publicly in 2006, and the later formulation that "the first Shopify store was our own" captured the process accurately precisely because the platform opportunity had not been the founders' original object of search (Shopify Inc., 2015).

At this stage Shopify was still primarily a software service for merchants, and the move from useful software to a broader platform depended on the decision to make parts of the system accessible to outside developers. On the company's third birthday in 2009, Shopify announced an API platform and the Shopify App Store, at a time when it served only a little more than 5,000 merchants. Developers could now create applications that accessed Shopify, add capabilities the core product did not contain and distribute those applications through a centralized store (Shopify, 2009). The architectural consequence was more important than the initially small catalogue of apps. Shopify no longer needed to predict and implement every function that heterogeneous merchants might eventually require; marketing, accounting, customer support, subscriptions, shipping, analytics and countless other capabilities could be provided by firms that Shopify did not own. This is precisely the tension captured later in the boundary-resources literature: APIs and related interfaces enable external contribution while simultaneously becoming instruments through which the Platform Owner defines access, protects the core and retains governance authority (Ghazawneh & Henfridsson, 2013). The original App Store contained fewer than a dozen applications, but by 2025 Shopify reported more than \$1.3 billion in payments to ecosystem developers and, by April 2026, more than 21,000 apps were available through the store (Shopify, 2018, 2026b). The system had become larger partly by becoming deliberately incomplete, with a stable core and interfaces through which others could create specialized complements.

Shopify was simultaneously bringing more of the merchant's operating environment into that core. In 2013 it launched Shopify Payments, reducing the need for merchants to arrange a separate payment gateway and merchant account, and introduced point-of-sale capabilities that carried the platform from online storefronts into physical retail (Shopify, 2013). Over the following years the company contin-

ued to extend the range of functions managed through the same administrative environment, while developers, payment providers, logistics services and other actors remained outside the corporate boundary. This mixture of centralization and external contribution is better understood as coevolution than as a simple movement from closed to open. Platform architecture, governance and the surrounding environment change together, as Tiwana, Konsynski and Bush emphasize, because a modification to one can alter incentives and dependencies elsewhere in the ecosystem (Tiwana et al., 2010). Shopify's merchant-facing platform therefore became neither a self-contained software product nor an uncontrolled network of third parties. It became an organized system in which the company governed important interfaces and transactions while relying on a large population of independent participants to create value that could not sensibly be produced inside the firm.

That architecture also helps explain why the rise of agentic commerce in 2026 should not be described simply as a threat to the storefront. Search engines and social platforms had already moved parts of product discovery away from merchant websites; generative AI extends the boundary because an agent need not merely display a link to a merchant's store. It can interpret an imprecise request, compare products offered by different merchants, continue the conversation when relevant information is missing and, increasingly, participate in the transaction itself, thereby shifting part of the commercial interaction into an interface that Shopify may enable without owning. Shopify described AI as the next major platform shift in commerce and, together with Google and other participants in commerce and payments, introduced the Universal Commerce Protocol, or UCP, as a common technical language through which agents, merchants and payment providers could coordinate discovery, carts, checkout, orders and post-purchase processes (Shopify, 2026a; Google, 2026a). UCP supports standardized commerce operations while allowing systems to negotiate capabilities and use different technical transports; Shopify Catalog, in turn, structures product information from large numbers of merchants so that external systems can search and interpret it, and its Catalog API makes that information accessible to developers (Shopify, 2026c; Google, 2026b; Shopify, 2026d).

The resulting strategic movement is almost the reverse of the one that produced Shopify in the first place. In 2004 Lütke needed a better store and therefore built software around the storefront; in 2026 Shopify was preparing for transactions in which the buyer might never visit that storefront at all. This does not imply that merchant websites, brands or direct customer relationships become irrelevant, because many purchases still depend on product presentation, service, loyalty and experiences controlled by the merchant. It means instead that the storefront can cease to be the obligatory entry point to every transaction. Shopify's response has been to make selected commerce capabilities usable beneath interfaces owned by others. In March 2026 the company announced Agentic Storefronts through which Shopify merchants could reach AI channels including ChatGPT, Microsoft Copilot, Google's AI Mode and Gemini from the Shopify Admin; it also opened parts of this

infrastructure to brands whose primary stores did not run on Shopify, allowing them to place products in Shopify Catalog and participate in agentic channels (Shopify, 2026e). In June it widened developer access to UCP and the Catalog API, presenting commerce increasingly as a capability that could be invoked inside many kinds of applications rather than as activity confined to a dedicated shopping destination (Shopify, 2026d,f).

This is where the Shopify case becomes especially useful for Platform Awareness, because the same organization occupies different positions depending on which platform is chosen as the focal system. Within the commerce ecosystem Shopify continues to exercise substantial Owner authority through its APIs, checkout, Catalog, payment infrastructure, App Store governance and the rules through which merchants, developers and partners interact. Within an agent-mediated commerce system, however, the role can be different. If Shopify supplies product data, checkout or payment capabilities that are required for the core transaction, the PBMC may locate it as a Provider; if its principal contribution in another configuration is standardization, integration or enabling coordination, Partner may be the more appropriate role. The platform literature may describe such an actor as a complementor, particularly where one platform increases the value of another, but *complementor* is not introduced here as a fifth PBMC perspective. The role has to be inferred from the Core Value Unit and transaction structure of the platform being analysed rather than from a permanent identity assigned to the company.

The distinction matters because Shopify does not own ChatGPT, Gemini or Copilot, and the fact that it does not own them does not by itself establish either strategic weakness or irrelevance. Through UCP and related interfaces Shopify can attempt to influence the standards by which commerce capabilities become legible to those systems, which may prove strategically more consequential than attempting to control every customer interface. This is a familiar problem in platform architecture: interfaces are not neutral connections, because the actor that defines or helps define them can influence which contributions are possible and under what conditions. Yet the present case should not be decided retrospectively while it is still unfolding. Agentic commerce remained uncertain in 2026; it was not clear which AI interfaces would persist, how often users would delegate purchases to agents, what governance and liability arrangements would emerge or how quickly standardized agentic checkout would spread across countries, merchants and product categories. Google itself continued to describe particular UCP-powered functions as limited in availability (Google, 2026c). Platform Awareness therefore does not require declaring a new technology the inevitable dominant layer. It requires noticing when a change in interfaces, actors or transactions is substantial enough that the organization's position may need to be reconsidered.

Read across Shopify's history, the recurring pattern is not a sequence of ever larger products but a series of changes in the system through which the same organization creates value. The merchant's own software became a service for other merchants; the service acquired APIs and an application ecosystem; commerce

functions such as payments and point of sale moved closer to the platform core; and, when external AI interfaces began to mediate product discovery and transactions, Shopify started adapting its infrastructure so that it could participate in systems it did not orchestrate. None of these movements requires the company to possess one immutable platform role. Shopify can be Owner of one platform while simultaneously acting as Provider or Partner in another, just as the platform it orchestrates can itself become a component of a wider inter-platform ecosystem. That possibility completes the historical movement developed through the preceding cases: Edward Lloyd occupied a valuable position in a maritime system he did not own, Western Union reused an existing network for another transaction, Sears scaled through external infrastructure, and Amazon recognized a platform inside its own organization. Shopify adds the reciprocal lesson that becoming a successful Platform Owner does not end the problem of orientation, because a platform can itself become a participant in another platform.

ACCOMPANYING CASE STUDY

The End of Online Stores? Shopify's Bet on AI Agents
Platform Generation

CHAPTER 8

Platform Awareness

Naming the pattern.

Platform Awareness is an organization's explicit and revisable understanding of the platforms in which it operates, the layer/position and role it occupies within each, and the strategic possibilities, dependencies and implications arising from those positions and their interconnection.

01

WHERE ARE WE?

Identify the relevant platforms and layers.

02

WHO ARE WE HERE?

Understand position and counterpart roles.

03

WHAT IS POSSIBLE?

Recognize opportunities, dependencies and implications.

Awareness comes before choice. The point is not to make every organization a Platform Owner; it is to make position visible before strategy begins.

By the time the historical sequence reaches Shopify, the distance between the cases is so great that any claim to have discovered one recurring platform architecture would be implausible. Edward Lloyd worked in a seventeenth-century coffee house whose commercial significance depended on face-to-face proximity, handwritten information and the institutional practices of maritime London; Western Union operated a nineteenth-century telegraph network; Sears combined catalogues, warehouses, railways and public postal infrastructure; Amazon's case turns on software services and distributed computing; and Shopify ends, for the moment, with machine-readable commerce protocols and AI-mediated interfaces. Their technologies, governance arrangements and transactions are not interchangeable. What makes them comparable is instead the way the strategic interpretation changes when the organization is located within a system larger than the product, asset or activity that initially dominates attention.

The recurring movement is therefore positional rather than developmental. The cases do not describe a progression from "non-platform" to "platform", nor do they imply that strategic maturity culminates in Platform Ownership. Lloyd's room becomes more intelligible when placed inside maritime information and exchange; Western Union's telegraph acquires another economic meaning when it coordinates money transfer; Sears's own assets explain only part of a business whose reach depended on infrastructures outside the firm; Amazon's internal services look different once their users need not be Amazon employees; and Shopify remains Owner of one platform while becoming a participant in other systems through which commerce may increasingly be accessed. This is consistent with research showing that firms can occupy several ecosystems and several platform relationships at once, but it shifts the analytical anchor from one selected platform to the organization whose strategically relevant positions may be distributed across several of them (Bosch-Sijtsema & Bosch, 2015; Shipilov & Gawer, 2020; Staykova & Damsgaard, 2016; Tavalaei & Cennamo, 2021; Schreieck et al., 2024; Carballa-Smichowski et al., 2026).

The definition deliberately treats awareness as a state of organizational understanding rather than assuming from the outset that it is a distinct organizational capability. An organization may possess a more or less explicit, coherent and revisable account of its platform positions; whether the ability to develop, maintain and update that account constitutes a measurable capability is a separate empirical question taken up in the research agenda. This distinction prevents the proposed construct from borrowing causal or performance implications that the historical cases cannot establish.

8.1 Reading the cases as connected positions

The figures in this section use a compact form of the Platform Business Model Canvas that will be formalized in Chapter 9. The Core Value Unit, or CVU, remains at the centre, with Owner above it, Provider to the left, Consumer to the right and

Partner below. The purpose here is not to replace the full PBMC or to pretend that every historical arrangement can be coded without ambiguity. It is to reread the cases with one stable coordinate system so that a change of reference platform becomes visible. The figures are consequently analytical reconstructions: they summarize the interpretation developed from the historical evidence, while the surrounding prose identifies where the reconstruction is more secure and where it remains deliberately provisional (Eisape, 2019, 2020a; Eisape, 2020b).

Edward Lloyd: Owner here, participant there

Edward Lloyd unquestionably controlled something important: his coffee house. He chose the premises, employed staff, gathered news, cultivated customers, hosted auctions and developed services around the room. The surviving sources show that the establishment became useful not only because it served coffee but because maritime information and commercially relevant people repeatedly converged there. The Hudson's Bay Company paid "Mr. Loyd the Coffeeman" for intelligence concerning its ships in 1693, while ships and other goods were sold through the premises; Lloyd briefly extended the information service beyond the room through *Lloyd's News* (Rich, 1958; Wright & Fayle, 1928; Lloyd's, 2013; Lloyd, 1696). Read narrowly, the coffee house can therefore be reconstructed as a platform-like system in which Lloyd occupied the Owner position because he controlled the place, access conditions and several of the services through which the focal value was created.

The same evidence becomes more interesting when the analytical boundary moves outward. Lloyd did not own the ships entering the Thames, the port, the post, London's merchant community or the marine-insurance practices that predated him; the historical record also does not justify making him the designer of the later Lloyd's insurance market (Wright & Fayle, 1928; Nature, 1928; Hodgson, 2014). Yet his establishment increasingly functioned as a useful node inside that wider maritime information and exchange environment. Captains and correspondents brought information, merchants and shipowners needed it, brokers and other intermediaries sought counterparties, and auctions added further transactions to the same location. Figure 8.1 therefore represents Lloyd as Owner in the first frame and as a Partner-like coordinating node in the second. The second frame is intentionally labelled as distributed in its governance: it is not a claim that seventeenth-century maritime London possessed one formally designed Platform Owner, but an analytical reconstruction of a wider system in which institutions, customs, brokers, merchants and information sources jointly shaped coordination.




Figure 8.1. Edward Lloyd as Owner in the coffee-house system and as a coordinating participant in a wider maritime information and exchange system. The second frame is an analytical reconstruction with distributed governance rather than a claim that maritime London possessed one formal Platform Owner.

The positional movement matters because it explains why Lloyd could become strategically consequential without owning the maritime system around him. By improving the information service, attracting relevant actors, hosting transactions and building reputation around an asset he controlled, he could increase the usefulness of his node to a much larger system. The later institutional history should not be read backwards as evidence that Lloyd planned the insurance market that inherited his name; it does show, however, how an Owner position in a comparatively small system can be connected to a different role in a larger one, and how that connection can alter the strategic meaning of the first position.

Western Union: same infrastructure, different transaction

Western Union produces a different configuration because the organization remains Owner while the focal transaction changes. By the late nineteenth century the company possessed telegraph lines, offices, operators, accounting routines and a trust structure capable of coordinating information across great distances. In the ordinary telegram business, the message itself was the focal value desired by the customer. Money transfer reorganized much of the same infrastructure around another outcome. A sender deposited money in one office; information concerning amount, recipient and payout was transmitted through the telegraph network; another office then paid the recipient from local funds. The exact banknotes placed on the first counter did not pass through the wire, while trusted information did (Western Union Telegraph Company, 1995).

Figure 8.2 makes the distinction explicit. Western Union remains Owner in both frames and relies on many of the same organizational capabilities, but the CVU is no longer the same. In the telegraph platform the communication is the focal value. In the money-transfer platform communication becomes an enabling mechanism for remote access to funds. That is why the case is analytically richer than a conventional statement that Western Union diversified into another product:

the same network, offices and trust arrangements support a different platform logic because the transaction organized through them has changed.




Figure 8.2. Western Union as Owner of two related platform configurations. Much of the infrastructure is shared, while the CVU and focal transaction change from transmitted message to remote access to funds.

The longer historical arc reinforces the point. Telegraphy eventually disappeared from the centre of Western Union’s business, while money transfer survived and became the company’s defining activity; modern corporate filings still date the consumer-to-consumer service to 1871 (Western Union Company, 2007). A platform position can therefore change economically even where a large part of the technical and organizational substrate remains recognizable.

Sears: Owner here, Consumer there

Richard Sears presents almost the inverse problem because the focal business depended on infrastructures outside the firm. Sears, Roebuck controlled its catalogue, assortment, advertising, order processing and increasingly sophisticated fulfilment machinery and can therefore be represented as Owner of the mail-order platform. National scale, however, depended on railway and postal systems whose governance remained elsewhere. Railways carried merchandise over distance; the postal system carried catalogues, orders and payments and increasingly reached rural households directly through Rural Free Delivery; Parcel Post later altered the practical economics of shipping heavier goods. The Smithsonian records that Sears mailed 2.3 million catalogues in 1910, while the launch of Parcel Post in 1913 was followed by a substantial increase in order volume (National Postal Museum, 2016; National Museum of American History, 2015; Pope, 2013; National Postal Museum, 2025).

A firm-centred diagram could place rail and post outside Sears as generic suppliers, but the evidence supports a more discriminating representation. The postal system can be analysed as a distinct coordinated system in which Sears appears as a user or Consumer, while the railway environment is better treated here as strategically consequential surrounding infrastructure rather than forced into a PBMC Core Frame. Nineteenth- and early twentieth-century rail transport involved multiple companies,

routes and governance arrangements, and the historical material used in this essay does not establish one single focal railway platform with a sufficiently specified Owner, CVU and role structure. Dependence is visible without being equated with weakness: both systems imposed conditions Sears could not determine, yet they also provided reach that would have been prohibitively expensive, and in the postal case institutionally impossible, for the company to reproduce independently. Figure 8.3 therefore makes both kinds of dependence visible while preserving the distinction between a coded platform and surrounding infrastructure.

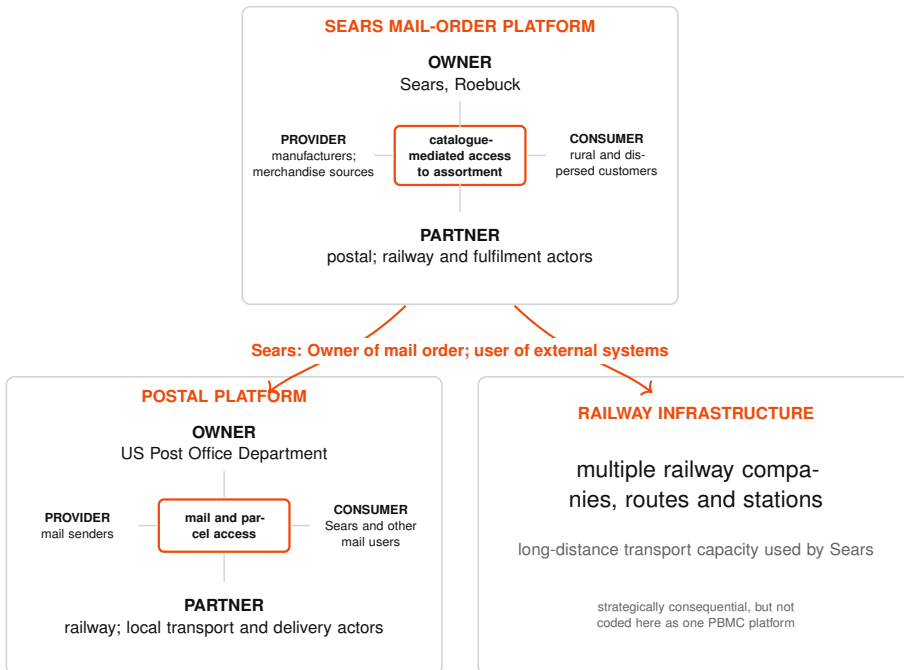


Figure 8.3. A Sears-centred Platform Landscape. Sears is represented as Owner of the mail-order platform and as a user of the postal platform. Railway transport is shown separately as surrounding infrastructure because the evidence used here does not justify collapsing multiple railway companies and governance arrangements into one PBMC Core Frame.

The biographical history makes this systems relationship unusually tangible without proving that Richard Sears conceptualized it in abstract terms. His first commercial opportunity emerged at a railway station; railway employees were among his earliest customers; advertising and orders travelled through the post; and later expansion benefited from public delivery infrastructure. What matters for Platform Awareness is not whether Sears possessed today's vocabulary, but that the business becomes difficult to understand if the analysis stops at the legal boundary of the company.

Amazon: discovering a platform inside the firm

Amazon reverses the direction again because the strategically relevant platform first appears inside the organization. Merchant.com exposed how tightly entangled Amazon’s technical capabilities had become when external retailers such as Target needed to use functions that appeared modular from the outside but were difficult to separate internally. Amazon consequently pushed toward clearer service boundaries and APIs, while internal product teams increasingly consumed reusable capabilities provided by other teams. Andy Jassy later described the earlier components as too “jumbled together”, and his retrospective accounts emphasize that the path toward AWS involved years of internal coordination problems, technical separation, developer feedback and repeated discussion rather than one moment of invention (Jassy, 2024, 2025).

Read through a PBMC, the internal service architecture begins to show differentiated Provider and Consumer roles around reusable capabilities before an external cloud business exists. Figure 8.4 is therefore deliberately an internal frame. Its significance lies in reversing the expected analytical sequence: management does not first identify an external platform and then use a PBMC to describe it; close analysis of repeated internal Provider–Consumer relationships, interfaces and reusable capabilities can itself reveal a platform-like structure that was absent from the initial strategic map.




Figure 8.4. Amazon before AWS externalization. Reusable technical capabilities can be read as an internal platform when differentiated Provider and Consumer roles, interfaces and recurring service relationships become visible.


Externalization then changes the frame without erasing its internal history. Developers and firms outside Amazon become Consumers, services receive explicit interfaces and prices, and capabilities originally built to solve Amazon’s own scaling problems become infrastructure for unknown applications. Werner Vogels’ arrival from Cornell in 2004 and the subsequent work on distributed systems helped translate these principles into operational architecture; recollections of early S3 design discussions over beer and bagels are memorable precisely because they sit beside engineering principles that treated component failure as normal and sought building blocks on which future services could depend (Cornell University, 2004; Vogels, 2022;

Amazon Staff, 2021). The platform was not literally complete before AWS existed, but the organizational relationships from which it could be recognized were already forming inside the firm.

Shopify: Owner here, Provider or Partner there

Shopify begins where many platform-strategy narratives want to end because it already owns and governs a substantial commerce platform. Merchants, developers and partners interact through interfaces and rules controlled to an important degree by Shopify, while payments, applications, point-of-sale capabilities and other services have moved closer to the platform core over time. Agentic commerce complicates the picture because product discovery and parts of the transaction may increasingly occur in interfaces owned by other organizations. The Universal Commerce Protocol and related catalogue and commerce APIs make this relationship concrete by creating technical mechanisms through which agents, merchants and payment providers can coordinate across discovery, checkout, order and post-purchase processes (Shopify, 2026a,c,d; Google, 2026a,b).

Figure 8.5 therefore places Shopify in two different frames. In the commerce platform it remains Owner. In an agent-mediated platform its position cannot be assigned from the company name alone. Where Shopify supplies catalogue, checkout, payment or merchant capabilities directly necessary for the focal transaction, Provider may be the more accurate PBMC role; where its contribution is primarily standardization, integration or enabling coordination, Partner may better describe the relationship. The question marks in the second frame are consequently intentional rather than unfinished design: the correct role follows from specification of the CVU and transaction, not from a permanent identity attached to Shopify.



Shopify: Owner here, Provider or Partner there

Figure 8.5. Shopify across two platform configurations. The question marks are analytical: Shopify's role in an agent-mediated commerce system depends on the focal CVU and transaction and may be coded as Provider or Partner rather than inferred from its Owner role in the commerce platform.

Recent research on inter-platform ecosystems makes the wider phenomenon explicit by showing configurations in which platforms direct users across platform

boundaries, aggregate or host other platforms and become complementors to one another without necessarily being governed by one overarching orchestrator (Carballa-Smichowski et al., 2026). The terminology differs from the PBMC: *complementor* remains useful when reporting that literature, but it is not introduced here as a fifth PBMC actor perspective. What the Shopify case adds to the historical sequence is the reciprocal lesson that becoming a successful Platform Owner does not resolve the problem of orientation, because the platform one governs can itself become a participant in another platform.

8.2 Roles are relational, and infrastructure can disappear from attention

Read together, the figures make it difficult to treat platform roles as permanent organizational identities. Ecosystem-as-structure research already distinguishes actors from the positions they occupy, and research on multiple ecosystems shows that the same firm can be central in one configuration and less influential, or differently engaged, in another (Adner, 2017; Bosch-Sijtsema & Bosch, 2015). In the terminology used here, Sears can therefore be Owner with respect to its mail-order platform and Consumer with respect to the postal platform while simultaneously depending on railway infrastructure that is not forced into a platform category; Shopify can be Owner of its commerce platform while appearing as Provider or Partner in a different PBMC; Lloyd can control the coffee-house system while participating in a wider maritime arrangement. The appropriate analytical question is not what an organization *is* in the abstract, but which role it occupies relative to a particular platform and CVU. That is the practical meaning of the second question on the Platform Awareness definition page: *Who are we here?*

This relational view also prevents the legal boundary of the firm from becoming the automatic boundary of analysis. Ecosystem research became important partly because value propositions often depend on legally independent actors whose activities have to align without being governed through ordinary hierarchy (Jacobides et al., 2018; Kretschmer et al., 2022). Edward Lloyd did not employ the captains, merchants and shipowners whose presence made his room useful; Sears did not employ the railway system or the US Post Office; Shopify does not own the external AI interfaces through which its commerce capabilities may increasingly be invoked. Yet stopping at the firm's legal boundary would remove systems without which the focal value proposition becomes weaker or changes altogether. The threshold is not simply externality or interdependence: as specified in the glossary, a platform requires a coordinating core, interface structure or equivalent organizing mechanism through which differentiated roles interact recurrently around a focal value-generating unit or transaction.

Infrastructure adds a further difficulty because successful infrastructure tends to become comparatively transparent in use. Star and Ruhleder emphasize that infrastructure is embedded in other structures and practices and often becomes especially visible when it fails, while Tilson, Lyytinen and Sørensen describe digital infrastructures as relational arrangements whose evolution involves continuing tensions of change and control (Star & Ruhleder, 1996; Tilson et al., 2010). The historical cases make the strategic implication concrete. Rail and post could become ordinary background conditions for Sears even though changes in those systems altered the reach and economics of mail order; Amazon's storage and compute capabilities could be treated as prerequisites of the "real" businesses precisely because teams had come to rely on them; Shopify merchants can use payments and applications without reconstructing every infrastructural relationship beneath the transaction. Platform Awareness therefore includes the effort to recover strategically consequential infrastructure from familiarity, while remaining selective enough not to reclassify every infrastructure, supplier or institution as a platform.

The cases also show that innovation can consist of repositioning an existing capability inside a different structure of interdependence. Western Union's wire remained a wire when the company organized money transfer; Amazon did not invent computing when it created AWS; Shopify's checkout does not cease to be checkout because another interface invokes it. What changes is the role of the capability, the Consumers to whom it is available, the CVU around which it is organized, the interface through which it can be accessed or the transaction it enables. A Platform Landscape is useful precisely because such changes can occur without the organization itself acquiring a new legal identity or building an entirely new technical substrate.

8.3 What the historical cases can establish

The historical material supports this conceptual argument only within clear limits. The cases were selected because different positional changes are unusually visible in them; they are neither a representative sample of organizations nor a basis for estimating the performance effects of Platform Awareness. Retrospective interpretation is especially vulnerable to hindsight because known successes encourage the reader to reconstruct earlier choices as more deliberate than the evidence allows. Edward Lloyd did not design the later Lloyd's insurance market, and the documentary record does not justify treating him as the visionary inventor of marine insurance. AWS likewise becomes less credible when reduced to a single moment of foresight, because participants themselves describe years of technical frustration, Merchant.com, internal services, developer feedback and repeated discussion (Barr, 2019; Miller, 2016). The method used throughout the essay therefore distinguishes documented historical behaviour from the platform interpretation imposed retrospectively on that behaviour.

What the cases can do is reveal mechanisms and positional movements that are sufficiently recurrent to motivate a conceptual question. A place controlled by Lloyd became more valuable as information and relevant actors converged around it; a communications network owned by Western Union supported a transaction different from the one for which the network had originally been built; Sears fitted its own commercial system to infrastructures it did not control; Amazon separated and externalized capabilities that had existed inside the firm; Shopify illustrates how an already successful platform can enter another platform's transaction as a participant. None of these observations proves that greater Platform Awareness causes superior performance. Together, however, they show why platform strategy can begin too late if it starts only after the focal platform and the organization's role within it have already been assumed, and they provide the empirical intuition from which the more systematic representational method in the next chapter is derived.

Locating the Organization

Platform Landscape and PBMC.

A concept concerned with orientation requires a representation capable of making orientation discussable, but the representation should not be confused with the phenomenon. Managers have long used maps and mapping metaphors precisely because strategic problems exceed what can be kept in one unaided mental picture. Fiol and Huff argued that broad managerial questions require a portfolio of different cognitive maps and that the interactions among those maps can be as important as the function of any one map (Fiol & Huff, 1992). A closer neighbouring approach is the Ecosystem Pie Model developed by Talmar and colleagues as a visual strategy tool for mapping, analysing and designing innovation ecosystems around complex value propositions (Talmar et al., 2020). Its existence is important for delimiting the contribution made here: visual ecosystem mapping itself is not the novelty claimed by the Platform Landscape. The narrower contribution proposed here is an organization-centred representation built from repeated PBMC structures that preserves specified CVUs and actor roles while making the same organization's changing positions across several platforms visible.

Platform architecture research provides a second foundation for that design. Baldwin and Woodard show that different platform traditions share architectural concerns with relatively stable cores, variable complementary components and interfaces, and discuss several representations that illuminate different aspects of platform architecture (Baldwin & Woodard, 2009). Ghazawneh and Henfridsson similarly demonstrate that boundary resources such as APIs are not merely technical connectors but governance mechanisms through which Platform Owners balance external contribution and control (Ghazawneh & Henfridsson, 2013). These literatures do not anticipate the Platform Landscape proposed here, but they support two design principles on which it depends: platforms may require several complementary representations, and interfaces are sufficiently consequential that a useful map must preserve the roles and transactions through which access, contribution and control are organized.

9.1 The PBMC Core Frame as a repeated coordinate system

The Platform Business Model Canvas was developed to make the multi-actor structure of platform business models explicit through actor perspectives, value propositions, transactions and governance (Eisape, 2019, 2020a; Eisape, 2020b). For the narrower task of locating several platforms at once, the full canvas contains more detail than a Landscape can carry legibly. This essay therefore uses an *adapted PBMC Core Frame*: a deliberately simplified representation that retains the Core Value Unit at the centre and the four PBMC actor perspectives around it—Owner above, Provider to the left, Consumer to the right and Partner below. The geometry should be understood as an adaptation for this essay rather than as a claim that the exact spatial arrangement is unchanged from every published version of the PBMC. Its purpose is to provide a repeated coordinate system: once the orientation is stable, the reader can compare several frames without relearning what each position means.




Figure 9.1. Adapted PBMC Core Frame used in this essay. The fixed orientation preserves the CVU and four PBMC actor perspectives as a repeated coordinate system for Landscape-level analysis.

The CVU remains central because the platform should not be defined merely by placing a familiar company name in the Owner position. The analytical object is the recurring value-generating structure in which actors occupy differentiated roles and interact around something that is created, exchanged, accessed or enabled. The Owner is the actor with substantial governance authority over that focal platform; the Provider contributes what is necessary for the CVU to be made available; the Consumer receives or uses the focal value; and the Partner contributes enabling capabilities, resources or coordination that materially support the platform without constituting the direct provision represented by the Provider role. These labels are analytical roles rather than permanent organizational types, so a coding decision is incomplete until the platform and CVU to which it refers have been specified.

Ambiguous cases, such as Shopify's position in agentic commerce, should be treated as prompts for closer analysis rather than resolved by company identity alone.

The Core Frame is intentionally incomplete. It does not replace the full PBMC, which remains the appropriate level for examining value propositions, transaction flows, governance mechanisms and the details through which a platform actually works. Its purpose is to preserve enough structure that several platforms can be placed in one Landscape without collapsing them into company logos and generic arrows. A Core Frame can consequently be wrong: a system initially labelled as a platform may lack meaningful differentiated roles or recurring transactions, a supposed Partner may prove to be a Provider once the CVU is specified, or an apparently peripheral API may turn out to govern access to a transaction that makes the platform strategically important.

9.2 A Platform Landscape as connected Core Frames

A Platform Landscape can now be defined more precisely as an organization-centred configuration of connected PBMC Core Frames. Each frame represents a platform relevant to the strategic question through its CVU and the Owner, Provider, Consumer and Partner perspectives; frames are connected where actors, capabilities, transactions or dependencies cross platform boundaries, allowing several positions of the focal organization—and, where analytically useful, other shared actors—to become visible at the same time. This is a proposition of the present essay rather than an established method imported from the literature. Its justification rests on observations that are already well supported: firms can participate in several ecosystems and take different positions within them, platforms can be deliberately integrated with other platforms, and platforms themselves can participate in wider inter-platform ecosystems (Bosch-Sijtsema & Bosch, 2015; Tavalaei & Cennamo, 2021; Schrieck et al., 2024; Carballa-Smichowski et al., 2026).

The simplest connection occurs when one legal entity appears in two frames. Sears can be represented as Owner of a mail-order platform while appearing as Consumer in the postal platform and depending separately on railway infrastructure; Shopify can remain Owner of its commerce platform while acting as Provider or Partner in an agent-mediated commerce platform. The connection between frames is therefore not the vague statement that two companies “have a relationship”. It identifies a role transition across two specified platform structures. Figure 9.2 abstracts the case-specific diagrams from Chapter 8 into a generic pattern.

Internal, focal and surrounding are positions in the Landscape

Once several Core Frames are connected, it becomes useful to describe their relation to the focal organization and strategic question without turning those labels into a




Figure 9.2. Connected PBMC Core Frames. The same organization can occupy different roles in different platforms; the connection makes the relational nature of those roles explicit.

taxonomy of platform types. An *internal* frame represents a platform or platform-like capability operating primarily within the organizational boundary, as Amazon’s internal services did before externalization. A *focal* frame is the platform currently at the centre of the strategic question, regardless of whether the organization owns it. A *surrounding* or inter-platform frame represents another platform on which the focal system depends, with which it interacts or within which it is itself positioned. The same AWS platform can therefore be surrounding from the perspective of a startup, focal from the perspective of an AWS strategy team and part of Amazon’s broader platform portfolio from the perspective of the corporation. What changes is the point of observation, not the ontology of AWS.

This relativity is also why a Landscape must be drawn for a strategic question rather than for the organization “in general”. A resilience inquiry may make cloud, identity, energy and logistics dependencies salient; an innovation question may instead foreground developers, standards and adjacent platforms. Fiol and Huff’s argument that managers need different maps for different strategic purposes is relevant because comprehensiveness is not the same as usefulness (Fiol & Huff, 1992). A Landscape that attempts to include every supplier, regulator, technology, institution and relationship can become less informative precisely because it no longer distinguishes what is consequential to the decision under examination. Standards, laws or institutions that are strategically important but are not convincingly platforms should therefore be annotated as context rather than forced into Core Frames merely to produce a visually uniform picture.

The organization around which the Landscape is built remains the analytical anchor, but it need not be the only actor linking frames. A shipowner may connect a shipping PBMC to an insurance PBMC; a payment company may connect merchant commerce to a card network; a developer may be Provider in one platform and Consumer of development infrastructure in another. Adner’s structural view of ecosystems is useful here because links relevant to a value proposition need not all

connect directly to the focal actor (Adner, 2017). The Landscape is consequently organization-centred without being organization-exclusive.

9.3 The Landscape and the full PBMC correct one another

The relationship between Landscape and PBMC is recursive rather than sequential. A first Landscape may reveal a platform that deserves close analysis; the full PBMC can then test the assumed roles, CVU, transactions and governance of that platform; the resulting analysis may confirm the initial frame, alter the organization's position within it or expose another platform that the first map omitted. Amazon illustrates movement from close analysis outward because repeated Provider–Consumer relationships among internal services can reveal a platform that had not initially appeared in the strategic map. Sears illustrates movement in the opposite direction because starting with the mail-order system pushes analysis outward when its reach cannot be understood adequately without the postal platform and surrounding railway infrastructure.




Figure 9.3. Analytical recursion between PBMC and Platform Landscape. Close analysis can confirm or revise the map and may expose additional platforms that require analysis in turn.

A map produced by this process is still not awareness. It is a representation of an interpretation and can be wrong even when it is visually elegant. Different organizational functions may construct different Landscapes because they possess different information, disagree about the CVU or assign different strategic significance to the same dependency. Distributed situation awareness suggests that such differences should not automatically be eliminated; they can reveal where organiza-

tional knowledge is fragmented and where locally reasonable views fail to produce a coherent account of the whole (Stanton et al., 2006). The value of the method lies in making those assumptions contestable. Why is this actor classified as Partner rather than Provider? Which transaction makes this system a platform rather than a loose supplier network? Who can change the interface or access rules? What part of the focal value proposition would disappear if the connection failed?

The PBMC therefore supplies the grammar from which the Landscape is constructed, while the Landscape supplies the context that prevents the PBMC from becoming too local. Without PBMC logic, a Landscape risks becoming a diagram of organizations and arrows whose analytical meaning depends on whoever drew it. Without a Landscape, a PBMC can describe one platform very well while leaving the organization unaware of the platforms immediately around it. Platform Awareness develops through the movement between the two: relevant systems become visible, their internal logic is tested closely, and the organization's several positions can then be interpreted together rather than one at a time.

From Awareness to Strategy

Awareness before choice.

The analytical sequence changes the beginning of platform strategy rather than dictating its conclusion. Once several positions are visible at the same time, “Should we build a platform?” becomes one possible strategic question among several. The organization may indeed have a credible opportunity to build and govern a new platform, but it may instead deepen a Provider role in a platform owned by somebody else, become a Partner whose contribution is difficult to substitute, externalize an internal platform, connect systems through a standard, diversify a dependency across several platforms or withdraw from a relationship whose governance has become too constraining. These alternatives are not stages in a ladder leading toward ownership, and the cases provide no basis for ranking them universally.

This distinction is consistent with research on dynamic capabilities, although Platform Awareness should not be collapsed into that broader framework. Teece’s formulation of sensing, seizing and reconfiguring concerns the organizational capacities through which firms recognize opportunities and threats, commit resources and transform themselves in changing environments (Teece, 2007). The *capacity to develop and revise* Platform Awareness has an obvious family resemblance to sensing because it concerns how an organization notices relevant systems, positions and changes before committing to a course of action. Platform Awareness itself, however, denotes the resulting state of explicit and revisable understanding, and its analytical object is narrower. A firm can be highly attentive to customer demand or technological change while misunderstanding which platform controls the interface through which the opportunity will be realized; conversely, recognizing that an internal capability is platform-like does not establish whether the firm should invest in externalizing it. Platform Awareness can inform the interpretation that precedes seizing, but it does not substitute for strategic choice.

10.1 Position before ownership

Ownership remains strategically important because Platform Owners can shape access, interfaces, governance and the conditions under which other actors create and capture value. Yet those rights come with corresponding demands. Helfat and Raubitschek argue that firms seeking to profit from innovation in digital platform ecosystems require dynamic and integrative capabilities because they must coordinate technological change and a population of participants that cannot simply be directed through hierarchy (Helfat & Raubitschek, 2018). Meta-organizational research makes the same difficulty visible from another angle: platform ecosystems involve governance among legally independent actors whose interests need not align with those of the focal firm (Kretschmer et al., 2022). The strategic value of an Owner position therefore depends on the attractiveness of the system, the governance rights ownership confers and the organization's ability to exercise those rights effectively.

The historical cases resist any simple assumption that greater control over more layers is necessarily superior. Sears benefited precisely because national rail and postal infrastructures already existed at scales the company could not plausibly reproduce; its dependency on them was real, but so was the reach they created. The analytical treatment need not be identical: the postal system can be coded as a focal platform for a specified transaction, while the more fragmented railway environment may remain surrounding infrastructure unless a particular railway platform is specified. Western Union's case shows that an organization can create a new platform logic by changing the transaction supported by infrastructure it already controls, while Amazon shows that a capability can become strategically different when its Consumers move beyond the firm. Shopify adds a further possibility because influence over an interface or standard can matter even when the company does not own the higher-level customer interface. The relevant strategic evaluation is therefore comparative and positional: which configuration offers an attractive combination of value creation, governance rights, dependence, substitutability and organizational capability for this organization in this system?

The Landscape helps make those comparisons more concrete. If Sears is represented only as a retailer, postal reform appears as environmental context; once the mail-order PBMC is connected to the postal platform and railway infrastructure is made explicit alongside it, the same reform becomes a change in the economics of a specific dependency. If Shopify is represented only as a Platform Owner, an AI interface can appear simply as a competitor or threat; once Shopify's commerce platform is shown as a participant in another transaction, questions of standards, Provider roles and interface governance become visible. The analytical gain lies less in producing an answer automatically than in changing which strategic questions can be asked before commitments are made.

10.2 Landscapes move

A Platform Landscape is necessarily a snapshot because architecture, governance and environmental conditions coevolve (Tiwana et al., 2010). Standards can reduce the bargaining power attached to one interface while increasing the reach of another; regulation can change access rights; an Owner can alter prices, ranking rules or technical conditions; an internal service can become externally consumable; a new intermediary can insert itself between an organization and its customers. Digital infrastructures intensify this dynamism because they are relational and can change while remaining deeply embedded in existing systems (Tilson et al., 2010). For that reason Platform Awareness should not be understood as possession of one definitive map. It is an explicit but revisable understanding of platform positions; the organizational capacity to recognize when that understanding has become outdated, and to revise it, is a related capability that remains to be tested empirically.

The Shopify case gives the issue contemporary urgency without making artificial intelligence the theoretical foundation of the argument. Agentic commerce remained unsettled in 2026, and neither the durability of particular interfaces nor the governance of autonomous transactions could be treated as established. The appropriate response to uncertainty is not to redraw the Landscape around every technological fashion. Instead, the analyst can ask whether the structure itself is changing: are different actors entering the transaction, has a new interface become a meaningful point of access or control, has a capability moved from internal to external use, have governance rights changed, or has another platform become necessary for delivery of the focal CVU? Where the answer is largely no, the new technology may be important without constituting a new platform position; where several of these relationships change together, the Landscape probably requires revision.

10.3 Limits and a Research Agenda

The conceptual ambition of the essay should remain proportionate to its evidence. The historical cases show that it is analytically productive to reconstruct organizations through several platform relationships, and the literature provides substantial adjacent foundations in situation awareness, distributed awareness, platform architecture, ecosystem strategy, multiple-ecosystem participation, infrastructure studies, dynamic capabilities and inter-platform ecosystems. The PBMC supplies an established framework for examining individual platform business models, while the connected Core Frame and Platform Landscape are proposed here as an organization-centred representational extension. What has not been demonstrated is that Platform Awareness is an empirically validated construct, that the organizational capacity to develop and revise it forms a statistically distinct capability, or that organizations displaying stronger Platform Awareness systematically perform better.

Those limitations suggest a research programme rather than a conclusion. Qualitative work could begin by asking managers from different functions within the same organization to construct Platform Landscapes independently and then comparing which platforms they identify, which roles they assign to the organization, how they specify CVUs and where they locate governance. Divergence among maps would not by itself indicate error; it could reveal the distribution of organizational knowledge and provide a basis for investigating which differences are consequential. Full PBMC analyses could then be used to test whether initial role assignments survive closer examination, whether supposed platforms actually display differentiated actor roles and recurring transactions, and whether important dependencies lie in systems that were absent from the first maps. Such work should not be restricted to retrospective success cases: negative or deviant cases in which organizations overlooked, misidentified or recognized too late a strategically consequential platform relationship would be especially important for separating a useful construct from a vocabulary that merely redescribes successful outcomes after the fact. Longitudinal designs could additionally examine whether Landscapes change before, during or only after consequential shifts in interfaces, governance or dependencies become visible in organizational decisions.

A later measurement programme would have to establish construct validity before attempting performance claims. Dimensions of Platform Awareness such as the identification of relevant platforms, correct localization of roles, recognition of interface and governance dependencies, and detection of emerging platform layers might form a coherent construct, but they might also collapse empirically into more general strategic sensing, ecosystem awareness or digital capability. A separate measurement question is whether the capacity to create and revise that understanding behaves as an organizational capability rather than merely reflecting individual expertise or existing routines. Discriminant validity is therefore not a technical afterthought; if Platform Awareness cannot be distinguished from established concepts, the new term would add little. Scenario-based measures may prove more informative than self-reported awareness because managers could be presented with unfamiliar cases, asked to identify and connect the relevant platforms, and compared with independently developed expert models before outcomes are known.

Only after such work would it become reasonable to test whether stronger Platform Awareness contributes to earlier recognition of changes in governance or interfaces, more effective diversification of critical dependencies, more frequent discovery of internal capabilities that can be externalized, or better decisions not to pursue Platform Ownership where the organization lacks a credible position. Those are empirical questions left open by the present essay.

The historical sequence can therefore end without turning its protagonists into practitioners of a framework they never knew. Edward Lloyd's importance does not depend on making him the designer of modern insurance; the case is revealing because a comparatively small asset he controlled became increasingly entangled with a wider system of information and exchange. Western Union demonstrates

that the transaction supported by a network can change even while much of the infrastructure remains. Sears shows that ownership of the focal business model can coexist with deep dependence on external platforms whose scale creates opportunity as well as constraint. Amazon reveals how a platform can become visible through close analysis of internal Provider–Consumer relationships before it is externalized. Shopify shows why the problem does not disappear once an organization becomes a successful Platform Owner, because the platform it governs can itself become a participant in another platform.

Taken together, the cases support a modest proposition with significant strategic consequences. Before an organization chooses where it wants to go in a platform economy, it benefits from knowing which systems already shape its choices, which positions it occupies within them and which relationships are changing. Platform strategy can then begin from a located organization rather than from an abstract ambition to become a platform.

Know where you are before deciding where to go.

Conceptual Glossary

Working definitions used in this essay.

Where concepts draw directly on or are adapted from established literature, their principal theoretical sources are indicated; concepts introduced as analytical propositions of this essay are identified as such.

ANALYTICAL LAYER

An observation level used to organize analysis: internal platforms, the focal platform or ecosystem, and surrounding or inter-platform systems. These are relative viewpoints determined by the strategic question, not mutually exclusive types of platform.

Analytical distinction used in this essay.

CORE VALUE UNIT (CVU)

The focal unit of value around which a platform's recurring transactions are organized and against which actor roles can be specified. Depending on the platform, the CVU may be created, exchanged, accessed or enabled and can take material, service, informational, financial or other forms.

Working definition used in the PBMC context. (Eisape, 2019, 2020a; Eisape, 2020b)

ECOSYSTEM

A set of interdependent actors and activities whose contributions jointly shape a value proposition, outcome or system. Ecosystem language is broader than platform language: an ecosystem need not be organized around a platform and may contain several platforms.

Synthesized from ecosystem-as-structure research. (Adner, 2017; Jacobides et al., 2018)

FOCAL PLATFORM / ECOSYSTEM

The system selected as the immediate object of analysis for a particular strategic question. "Focal" is relational: the same organization or system can be focal in one analysis and surrounding in another.

Working definition in this essay; informed by ecosystem-as-structure research. (Adner, 2017)

GOVERNANCE

The rules and mechanisms that shape participation, access, behaviour, rights, responsibilities and change within a system. Governance may operate through ownership and contracts, but also through standards, protocols, institutions, reputation, shared norms and technical interfaces.

Working definition in this essay; informed by platform and ecosystem research. (Gawer & Cusumano, 2014; Jacobides et al., 2018; Ghazawneh & Henfridsson, 2013)

INFRASTRUCTURE

A relational and embedded set of material, technical, organizational or institutional arrangements that supports recurring activity and may become comparatively transparent in use. Infrastructure is not synonymous with platform, although infrastructure can contain or become part of platform structures.

Working definition informed by infrastructure studies. (Star & Ruhleder, 1996; Tilson et al., 2010)

INTERFACE

A point at which actors, components or systems connect and interact. Interfaces may be technical, such as APIs, but can also be organizational, contractual or institutional; in platform settings they may simultaneously enable participation and embody governance choices.

Informed by platform-architecture and boundary-resource research. (Baldwin & Woodard, 2009; Ghazawneh & Henfridsson, 2013)

INTERNAL PLATFORM

A platform or platform-like capability located primarily within the organizational boundary and relevant to the strategic question. Not every shared internal technology is a platform; the concept becomes analytically useful where differentiated actor roles, recurring interactions and coordinating interfaces or governance are present.

Adapted from internal/company-specific platform research, especially Gawer (2014). (Gawer & Cusumano, 2014)

NETWORK

A pattern of connections or relationships among actors or nodes. A network establishes that actors are connected but does not by itself establish a platform architecture, recurring multi-actor transactions or platform governance.

Working distinction used in this essay.

NETWORK EFFECT

A change in the value of a product or service associated with changes in the number or composition of users. Network effects are important to platform economics, but their presence alone does not make every network a platform.

Based on foundational network-effects research. (Rohlf's, 1974; Katz & Shapiro, 1985)

PLATFORM

A structured arrangement that enables and coordinates recurring interactions or transactions among distinct actor roles through interfaces and governance mechanisms. Mere interdependence is insufficient: the arrangement must contain an identifiable coordinating core, interface structure or equivalent organizing mechanism through which differentiated roles interact recurrently around a focal value-generating unit or transaction. A platform does not have to be digital and need not always be governed by one identifiable owner.

Working synthesis for this essay from platform-economics and platform-strategy research.

(Rochet & Tirole, 2003; Rysman, 2009; Gawer & Cusumano, 2014; Cusumano et al., 2019)

PLATFORM AWARE- NESS

An organization's explicit and revisable understanding of the platforms in which it operates, the layer/position and role it occupies within each, and the strategic possibilities, dependencies and implications arising from those positions and their interconnection. The term denotes a state of organizational understanding; whether the capacity to develop, maintain and revise that understanding constitutes a distinct organizational capability is left as an empirical question.

Proposed in this essay; conceptually related, but not equivalent, to situation awareness, distributed situation awareness, situational platform awareness and media-education conceptions of platform awareness. (Endsley, 1995; Stanton et al., 2006; Burström et al., 2016; Spurava, 2024)

PLATFORM BUSI- NESS MODEL CANVAS (PBMC)

An analytical framework for examining a focal platform through its actor perspectives, Core Value Unit, value propositions, transactions and governance. In this essay the full PBMC tests and deepens the simplified representations used in the Platform Landscape.

Based on the Platform Business Model Canvas literature. (Eisape, 2019, 2020a; Eisape, 2020b)

PBMC ACTOR ROLES

Four relational coding positions used around a specified platform and CVU. The Owner exercises substantial governance authority; the Provider contributes what is directly necessary to make the CVU available; the Consumer receives or uses the focal value; and the Partner contributes enabling capabilities, resources or coordination that materially support the platform without constituting the direct provision represented by the Provider role. One organization can occupy different roles in different platforms.

Coding convention used in this essay; based on the PBMC actor grammar. (Eisape, 2019, 2020a; Eisape, 2020b)

PBMC CORE FRAME

An adapted and deliberately simplified PBMC representation used in this essay for Landscape-level analysis. It retains the CVU and four actor perspectives—Owner, Provider, Consumer and Partner—in a fixed geometry so that several platforms can be connected and compared without reproducing the full canvas.

Proposed as an adaptation in this essay; based on the PBMC actor grammar. (Eisape, 2019, 2020a; Eisape, 2020b)

PLATFORM ECOSYSTEM

The constellation of actors, complements, dependencies and governance relationships organized around or through a focal platform. A platform ecosystem is more specific than an ecosystem in general because the platform provides an important organizing structure.

Synthesized from platform and ecosystem research. (Gawer & Cusumano, 2014; Jacobides et al., 2018)

PLATFORM LANDSCAPE

An organization-centred configuration of connected PBMC Core Frames representing the platforms relevant to a strategic question. Frames are linked where actors, capabilities, transactions or dependencies cross platform boundaries, making the organization's simultaneous positions across several platforms visible. A Landscape is a provisional analytical representation rather than a permanent or exhaustive map.

Proposed in this essay; informed by multi-ecosystem, inter-platform, platform-architecture, ecosystem-mapping and strategic-mapping research. (Bosch-Sijtsema & Bosch, 2015; Shipilov & Gawer, 2020; Schrieck et al., 2024; Carballa-Smichowski et al., 2026; Talmar et al., 2020; Fiol & Huff, 1992; Baldwin & Woodard, 2009)

PLATFORM LOGIC

The coordination logic produced by structured roles, interfaces, recurring interactions or transactions, and governance. The term is deliberately technology-neutral: platform logic can appear in digital and non-digital settings.

Working definition in this essay.

PLATFORM-LIKE CO-ORDINATION LOGIC

A deliberately cautious term for historical or adjacent arrangements that exhibit some features associated with platform logic without claiming that they satisfy a contemporary platform definition. The distinction is particularly important in the deep-history discussion.

Term used in this essay for interpretive caution.

PLATFORM OWNER

An actor that exercises substantial governance authority over a focal platform. Platform Ownership is relational rather than an intrinsic identity: an organization can own one platform while participating in another.

Working definition in this essay; informed by platform-leadership research. (Gawer & Cusumano, 2002, 2014)

PLATFORM ROLE

An actor's position relative to a particular focal platform and its CVU. Roles are not fixed attributes of organizations; the same organization can occupy different roles across different platforms.

Relational working definition in this essay; supported by research on multiple-ecosystem participation. (Bosch-Sijtsema & Bosch, 2015; Tavalaei & Cennamo, 2021)

PLATFORM STRATEGY

Strategic choices about how an organization acts within, through, across or in relation to platforms. In this essay Platform Strategy is explicitly not synonymous with building or owning a platform.

Working definition in this essay.

SURROUNDING / INTER-PLATFORM SYSTEM

A platform, infrastructure or other structured system on which the focal platform depends, with which it interacts or within which it is itself positioned. The category is relative to the strategic question and does not imply that every surrounding institution, standard or technology is a platform.

Analytical category used in this essay; informed by multi-platform and inter-platform research. (Schrieck et al., 2024; Carballa-Smithowski et al., 2026)

TRANSACTION

A recurring exchange or transfer between actors. Depending on the platform, what moves may include goods, services, money, information, attention, rights, risk or access. In the PBMC, transactions make interdependence concrete.

Working definition used in the PBMC context. (Eisape, 2019, 2020a; Eisape, 2020b)

TRANSACTION, INNOVATION AND HYBRID PLATFORMS

A transaction platform primarily enables interaction or exchange between distinct participant groups. An innovation platform provides a technological foundation on which others develop complementary innovations. A hybrid platform combines both logics. This distinction is also used in the market-capitalisation figure.

Based on the transaction, innovation and hybrid-platform typology. (Cusumano et al., 2019)

References

- Abbo, S. & Gopher, A. (2020) 'Plant domestication in the Neolithic Near East: The humans-plants liaison', *Quaternary Science Reviews*, 242, 106412. DOI 10.1016/j.quascirev.2020.106412.
- Addison, J. (1711) 'The Spectator, No. 46', *The Spectator*, 23 April 1711.
- Adner, R. (2012) *The Wide Lens: A New Strategy for Innovation*. New York: Portfolio/Penguin.
- Adner, R. (2017) 'Ecosystem as structure: An actionable construct for strategy', *Journal of Management*, 43(1), pp. 39–58. DOI 10.1177/0149206316678451.
- Amazon Staff (2021) 'The deceptively simple origins of AWS', About Amazon. Seattle, WA: Amazon.com, Inc.
- Amazon Web Services (2006b) 'Announcing Amazon Elastic Compute Cloud (Amazon EC2) – beta', 24 August. Seattle, WA: Amazon Web Services.
- Amazon.com, Inc. (2008) Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2007. Seattle, WA: Amazon.com, Inc.
- Armstrong, M. (2006) 'Competition in two-sided markets', *RAND Journal of Economics*, 37(3), pp. 668–691. DOI 10.1111/j.1756-2171.2006.tb00037.x.
- Baldwin, C.Y. & Woodard, C.J. (2009) 'The Architecture of Platforms: A Unified View', in Gawer, A. (ed.) *Platforms, Markets and Innovation*, pp. 19–44. Edward Elgar. DOI 10.4337/9781849803311.00008.
- Barr, J. (2006a) 'Amazon S3', AWS News Blog, 14 March.
- Barr, J. (2006b) 'Amazon EC2 Beta', AWS News Blog, 25 August.
- Barr, J. (2019) '15 Years of AWS Blogging!', AWS News Blog, 11 November.
- Besen, S.M. & Farrell, J. (1994) 'Choosing how to compete: Strategies and tactics in standardization', *Journal of Economic Perspectives*, 8(2), pp. 117–131. DOI 10.1257/jep.8.2.117.
- Bezos, J. (2007) '2006 Letter to Shareholders', Amazon.com, Inc., filed with the U.S. Securities and Exchange Commission.

- Bosch-Sijtsema, P.M. & Bosch, J. (2015) 'Plays nice with others? Multiple ecosystems, various roles and divergent engagement models', *Technology Analysis & Strategic Management*, 27(8), pp. 960–974. DOI 10.1080/09537325.2015.1038231.
- Burström, T., Grönlund, M. & Ranti, T. (2016) *The Role of High-impact SMEs in Finland*. Tekes Review 328/2016. Helsinki: Tekes – Finnish Funding Agency for Innovation. ISBN 978-952-457-616-1.
- Caillaud, B. (2019) *Curriculum Vitae*. Paris: Paris School of Economics.
- Caillaud, B. & Jullien, B. (2001) 'Chicken and egg: Competing matchmakers', CEPR Discussion Paper No. 2885. London: Centre for Economic Policy Research.
- Carballa-Smichowski, B., Duch-Brown, N., Martens, B. & Gomez-Losada, Á. (2026) 'Inter-platform ecosystems', *Strategic Management Journal*, 47(7), pp. 1840–1877. DOI 10.1002/smj.70070.
- Charles II (1675) *By the King. A Proclamation for the Suppression of Coffee-Houses*. London: John Bill and Christopher Barker, 29 December 1675.
- Chicago Tribune (1907) 'Richard W. Sears', 16 June. Reproduced in *Chicagology*, 'Richard W. Sears'.
- Cornell University (2004) *Computing and Information Science Annual Report 2004*. Ithaca, NY: Cornell University.
- Cusumano, M.A., Gawer, A. & Yoffie, D.B. (2019) *The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, and Power*. New York: Harper Business.
- Dietrich, L., Meister, J., Dietrich, O., Notroff, J., Kiep, J., Heeb, J., Beuger, A. & Schütt, B. (2019) 'Cereal processing at Early Neolithic Göbekli Tepe, southeastern Turkey', *PLOS ONE*, 14(5), e0215214. DOI 10.1371/journal.pone.0215214.
- Dietrich, O. (2016) 'Göbekli Tepe – The first 20 Years of Research', *Tepe Telegrams*, 2 June. German Archaeological Institute.
- Eisape, D. (2019) 'The Platform Business Model Canvas: A proposition in a design science approach', *American Journal of Management Science and Engineering*, 4(6), pp. 91–107. DOI 10.11648/j.ajmse.20190406.12.
- Eisape, D. (2020a) 'Comparing Platform Business Models: A Balanced Scorecard Approach based on the Platform Business Model Canvas', *Nordic Journal of Media Management*, 1(3), pp. 401–432. DOI 10.5278/njmm.2597-0445.5842.
- Eisape, D.A. (2020b) 'The Platform Business Model Canvas – Designing and Visualizing Platform Business Models', *International Journal of Innovative Studies in Sciences and Engineering Technology*, 6(2).

- Emmet, B. & Jeuck, J.E. (1950) *Catalogues and Counters: A History of Sears, Roebuck and Company*. Chicago: University of Chicago Press.
- Endsley, M.R. (1995) 'Toward a theory of situation awareness in dynamic systems', *Human Factors*, 37(1), pp. 32–64. DOI 10.1518/001872095779049543.
- Evans, D.S. (2003) 'The antitrust economics of multi-sided platform markets', *Yale Journal on Regulation*, 20(2), pp. 325–382.
- Ewens, M., Nanda, R. & Rhodes-Kropf, M. (2018) 'Cost of experimentation and the evolution of venture capital', *Journal of Financial Economics*, 128(3), pp. 422–442. DOI 10.1016/j.jfineco.2018.03.001.
- Fiol, C.M. & Huff, A.S. (1992) 'Maps for Managers: Where Are We? Where Do We Go from Here?', *Journal of Management Studies*, 29(3), pp. 267–285. DOI 10.1111/j.1467-6486.1992.tb00665.x.
- Gawer, A. (2000) *The Organization of Platform Leadership: An Empirical Investigation of Intel's Management Processes Aimed at Fostering Complementary Innovation by Third Parties*. PhD thesis. Cambridge, MA: MIT Sloan School of Management.
- Gawer, A. & Cusumano, M.A. (2002) *Platform Leadership: How Intel, Microsoft, and Cisco Drive Industry Innovation*. Boston, MA: Harvard Business School Press.
- Gawer, A. & Cusumano, M.A. (2014) 'Industry platforms and ecosystem innovation', *Journal of Product Innovation Management*, 31(3), pp. 417–433. DOI 10.1111/jpim.12105.
- Ghazawneh, A. & Henfridsson, O. (2013) 'Balancing platform control and external contribution in third-party development: the boundary resources model', *Information Systems Journal*, 23(2), pp. 173–192. DOI 10.1111/j.1365-2575.2012.00406.x.
- Google (2026a) 'New tech and tools for retailers to succeed in an agentic shopping era', Google Blog, 11 January 2026.
- Google (2026b) Handa, A. & Gupta, A., 'Under the Hood: Universal Commerce Protocol (UCP)', Google Developers Blog, 11 January 2026.
- Google (2026c) 'About the Universal Commerce Protocol (UCP) and UCP-powered checkout feature on Google', Google Merchant Center Help. Accessed 17 September 2026.
- Groman-Yaroslavski, I., Weiss, E. & Nadel, D. (2016) 'Composite sickles and cereal harvesting methods at 23,000-years-old Ohalo II, Israel', *PLOS ONE*, 11(11), e0167151. DOI 10.1371/journal.pone.0167151.
- Harari, Y.N. (2015) *Sapiens: A Brief History of Humankind*. New York: Harper.

- Helfat, C.E. & Raubitschek, R.S. (2018) 'Dynamic and integrative capabilities for profiting from innovation in digital platform-based ecosystems', *Research Policy*, 47(8), pp. 1391–1399. DOI 10.1016/j.respol.2018.01.019.
- Hershkovitz, I., Speirs, M.S., Frayer, D., Nadel, D., Wish-Baratz, S. & Arensburg, B. (1995) 'Ohalo II H2: A 19,000-year-old skeleton from a water-logged site at the Sea of Galilee, Israel', *American Journal of Physical Anthropology*, 96(3), pp. 215–234. DOI 10.1002/ajpa.1330960302.
- Historic American Buildings Survey (1933) Sears Roebuck & Company Mail Order Plant, Merchandise Building, 924 South Homan Avenue, Chicago, Cook County, IL. Survey HABS IL-1187-A. Washington, DC: Library of Congress.
- Hodgson, R. (2014) *Lloyd's: Law and Practice*. Abingdon: Informa Law/Routledge.
- Jacobides, M.G., Cennamo, C. & Gawer, A. (2018) 'Towards a theory of ecosystems', *Strategic Management Journal*, 39(8), pp. 2255–2276. DOI 10.1002/smj.2904.
- Jassy, A. (2022) '2022 Letter to Shareholders', About Amazon. Seattle, WA: Amazon.com, Inc.
- Jassy, A. (2024) '2023 Letter to Shareholders', Amazon.com, Inc., filed with the U.S. Securities and Exchange Commission, April 2024.
- Jassy, A. (2025) '2025 Letter to Shareholders', About Amazon. Seattle, WA: Amazon.com, Inc.
- Jullien, B. (2025) *Curriculum Vitae*. Toulouse: Toulouse School of Economics.
- Katz, M.L. & Shapiro, C. (1985) 'Network externalities, competition, and compatibility', *American Economic Review*, 75(3), pp. 424–440.
- Katz, M.L. & Shapiro, C. (1994) 'Systems competition and network effects', *Journal of Economic Perspectives*, 8(2), pp. 93–115. DOI 10.1257/jep.8.2.93.
- Kretschmer, T., Leiponen, A., Schilling, M. & Vasudeva, G. (2022) 'Platform ecosystems as meta-organizations: Implications for platform strategies', *Strategic Management Journal*, 43(3), pp. 405–424. DOI 10.1002/smj.3250.
- Lee, S. (1893) 'Lloyd, Edward', in Lee, S. (ed.) *Dictionary of National Biography*, Vol. 33. London: Smith, Elder & Co.
- Library of Congress (1875a) General Operating Department, Western Union Telegraph Building, New York. Photograph. Prints and Photographs Division, Library of Congress, Washington, DC. Reproduction no. LC-DIG-ppmsca-55661.
- Library of Congress (1875b) Bill, C.K., Western Union Telegraph Building, New York – 'Switch', General Operating Department. Stereograph. Prints and Photographs Division, Library of Congress, Washington, DC. LCCN 93508150.

- Library of Congress (2026) 'Sears, Roebuck, & Co. Goes Public', This Month in Business History. Washington, DC: Library of Congress.
- Liebowitz, S.J. & Margolis, S.E. (1994) 'Network externality: An uncommon tragedy', *Journal of Economic Perspectives*, 8(2), pp. 133–150. DOI 10.1257/jep.8.2.133.
- Lloyd, E. (1696) *Lloyd's News*, 15 October 1696. London.
- Lloyd's (2013) 'Lloyd's at 325: the story of Edward Lloyd', 15 February. London: Lloyd's. Archived official history article.
- Lloyd's (2020) 'Our full statement on the Lloyd's market's role in the slave trade', 18 June. London: Lloyd's.
- Lütke, T. (2010) 'Q&A with Tobias Lütke of Shopify', *Signal v. Noise*.
- Miller, R. (2016) 'How AWS came to be', *TechCrunch*, 2 July.
- Nadel, D. (1990) 'Ohalo II: A preliminary report', *Mitekufat Haeven*, 23.
- Nadel, D. (2006) 'Ohalo I – Shaldag Beach: A final report on an Epipalaeolithic–Neolithic workshop site in the Sea of Galilee', *Mitekufat Haeven*, 36, pp. 39–97.
- Nadel, D. (2017) 'Ohalo II: A 23,000-year-old fisher-hunter-gatherers' camp on the shore of fluctuating Lake Kinneret (Sea of Galilee)', in Enzel, Y. & Bar-Yosef, O. (eds.) *Quaternary of the Levant*. Cambridge: Cambridge University Press.
- Nadel, D., Weiss, E., Simchoni, O., Tsatskin, A., Danin, A. & Kislev, M. (2004) 'Stone Age hut in Israel yields world's oldest evidence of bedding', *Proceedings of the National Academy of Sciences*, 101(17), pp. 6821–6826.
- National Museum of American History (2015) 'Let your fingers do the shopping', *Object Project*. Washington, DC: Smithsonian Institution.
- National Park Service (2024) 'The Transcontinental Telegraph'. U.S. National Park Service. Updated 7 May 2024.
- National Postal Museum (2016) 'The Growth of the Mail', in *America's Mailing Industry* [virtual exhibition]. Washington, DC: Smithsonian Institution.
- National Postal Museum (2025) 'The 20th Century: Parcel Post'. Washington, DC: Smithsonian Institution.
- Nature (1928) 'A History of Lloyd's: From the Founding of Lloyd's Coffee-house to the Present Day', *Nature*, 122, pp. 267–268. DOI 10.1038/122267a0.
- Nobel Prize Outreach (2014) 'The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2014: Jean Tirole'. Stockholm: Nobel Prize Outreach.

- Odling-Smee, F.J. (1988) 'Niche-constructing phenotypes', in Plotkin, H.C. (ed.) *The Role of Behavior in Evolution*. Cambridge, MA: MIT Press, pp. 73–132.
- Odling-Smee, F.J., Laland, K.N. & Feldman, M.W. (2003) *Niche Construction: The Neglected Process in Evolution*. Princeton, NJ: Princeton University Press.
- Oldham, J. (2007) 'Insurance litigation involving the Zong and other British slave ships, 1780–1807', *The Journal of Legal History*, 28(3), pp. 299–318. DOI 10.1080/01440360701698437.
- Orton, W. (1870) *Government Telegraphs: Argument of William Orton, President of the Western Union Telegraph Company, on the Bill to Establish Postal Telegraph Lines, Delivered Before the Select Committee of the United States House of Representatives*. New York: Russells' American Steam Printing House.
- Ottawa Business Journal (2017) 'Entrepreneurial lessons: Shopify CEO Tobi Lütke on his first sale'. Ottawa.
- Palmer, S. (2007) 'Lloyd, Edward (c.1648–1713)', *Oxford Dictionary of National Biography*. Oxford: Oxford University Press. DOI 10.1093/ref:odnb/16829.
- Parker, G.G. & Van Alstyne, M.W. (2000) 'Information complements, substitutes, and strategic product design', *Proceedings of the Twenty-First International Conference on Information Systems (ICIS 2000)*, Brisbane, Australia, 10–13 December, Paper 2.
- Parker, G.G., Van Alstyne, M.W. & Choudary, S.P. (2016) *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*. New York: W.W. Norton & Company.
- Pinkham, C. (2011) interview in 'Chris Pinkham: veteran of the virtual', TechCentral.
- Pope, N.A. (2013) '100 Years of Parcels, Packages, and Packets, Oh My!', *National Postal Museum research article*. Washington, DC: Smithsonian Institution.
- PricewaterhouseCoopers (PwC) (2014) *Global Top 100 Companies by Market Capitalisation*. Market values as at 31 March 2014, with 2009 comparatives. Source data: Bloomberg and PwC analysis.
- PricewaterhouseCoopers (PwC) (2016) *Global Top 100 Companies by Market Capitalisation*. Market values as at 31 March 2016. Source data: Bloomberg and PwC analysis.
- PricewaterhouseCoopers (PwC) (2019) *Global Top 100 Companies by Market Capitalisation*. Market values as at 31 March 2019. Source data: Bloomberg and PwC analysis.
- PricewaterhouseCoopers (PwC) (2021) *Global Top 100 Companies by Market Capitalisation*. Includes market capitalisation as at 31 March 2020 and 31 March 2021. Source data: Bloomberg and PwC analysis.

- PricewaterhouseCoopers (PwC) (2026) Global Top 100 Companies by Market Capitalisation 2026. Market values as at 31 March 2026; source data reported as S&P Global Market Intelligence with PwC analysis.
- Printers' Ink (1914) 'Richard W. Sears', 8 October. Reproduced in *Chicagology*, 'Richard W. Sears'.
- Reid, J.D. (1879) *The Telegraph in America: Its Founders, Promoters, and Noted Men*. New York: Derby Brothers.
- Rich, E.E. (1958) *The History of the Hudson's Bay Company, 1670–1870. Volume I: 1670–1763. Publications of the Hudson's Bay Record Society, Vol. XXI*. London: Hudson's Bay Record Society.
- Rietveld, J. & Schilling, M.A. (2021) 'Platform Competition: A Systematic and Interdisciplinary Review of the Literature', *Journal of Management*, 47(6), pp. 1528–1563. DOI 10.1177/0149206320969791.
- Rochet, J.-C. & Tirole, J. (2003) 'Platform competition in two-sided markets', *Journal of the European Economic Association*, 1(4), pp. 990–1029. DOI 10.1162/154247603322493212.
- Rochet, J.-C. & Tirole, J. (2006) 'Two-sided markets: A progress report', *RAND Journal of Economics*, 37(3), pp. 645–667. DOI 10.1111/j.1756-2171.2006.tb00036.x.
- Rochet, J.-C. & Tirole, J. (2023) 'Jean-Charles Rochet and Jean Tirole introduce their article, Platform Competition in Two-Sided Markets (2003)', in Wacziarg, R. et al., 'Reflecting on the first 20 years of the Journal of the European Economic Association', *Journal of the European Economic Association*, 21(5), pp. e1–e33. DOI 10.1093/jeea/jvad045.
- Rohlf, J. (1974) 'A theory of interdependent demand for a communications service', *Bell Journal of Economics and Management Science*, 5(1), pp. 16–37. DOI 10.2307/3003090.
- Royal Museums Greenwich (2023) 'How coffee has shaped Britain's maritime history'. Greenwich: Royal Museums Greenwich.
- Rysman, M. (2009) 'The economics of two-sided markets', *Journal of Economic Perspectives*, 23(3), pp. 125–143. DOI 10.1257/jep.23.3.125.
- Schrieck, M., Ondrus, J., Wiesche, M. & Krcmar, H. (2024) 'A Typology of Multi-Platform Integration Strategies', *Information Systems Journal*, 34(3), pp. 828–853. DOI 10.1111/isj.12450.
- Shipilov, A. & Gawer, A. (2020) 'Integrating Research on Interorganizational Networks and Ecosystems', *Academy of Management Annals*, 14(1), pp. 92–121. DOI 10.5465/annals.2018.0121.

- Shopify (2009) 'Shopify launches API platform and App Store', 2 June. Ottawa: Shopify.
- Shopify (2013a) 'Shopify launches Shopify Payments', 12 August. Ottawa: Shopify.
- Shopify (2018a) 'Shopify launches all new App Store', 5 September. Ottawa: Shopify.
- Shopify (2024) 'Tobi Lütke: Shopify founder's success advice', 5 December.
- Shopify (2026a) 'The agentic commerce platform: Shopify connects any merchant to every AI conversation', 11 January.
- Shopify (2026b) 'The developers behind Shopify's \$1.3 billion app ecosystem', 28 April.
- Shopify (2026c) Universal Commerce Protocol. Technical overview and specification. Accessed 17 September 2026.
- Shopify (2026d) 'Agentic commerce for every developer: The Spring '26 Edition', 17 June.
- Shopify (2026e) 'Millions of merchants can sell in AI chats', 24 March.
- Shopify (2026f) 'Spring '26 Edition: Five apps that show what Catalog API and UCP make possible', 17 June.
- Shopify Inc. (2015) Form F-1 Registration Statement. Filed with the U.S. Securities and Exchange Commission, 14 April 2015, File No. 333-203295.
- Shopify Inc. (2026) Annual Report and Form 10-K/A for the Fiscal Year Ended December 31, 2025. Ottawa: Shopify Inc., filed with the U.S. Securities and Exchange Commission.
- Snir, A., Nadel, D., Groman-Yaroslavski, I., Melamed, Y., Sternberg, M., Bar-Yosef, O. & Weiss, E. (2015) 'The origin of cultivation and proto-weeds, long before Neolithic farming', *PLOS ONE*, 10(7), e0131422. DOI 10.1371/journal.pone.0131422.
- Spengler, R.N. III (2020) 'Anthropogenic seed dispersal: Rethinking the origins of plant domestication', *Trends in Plant Science*, 25(4), pp. 340–348. DOI 10.1016/j.tplants.2020.01.005.
- Spurava, G. (2024) Towards Platform Awareness in Media Education. Tampere University Dissertations, Vol. 1122. Tampere: Tampere University. ISBN 978-952-03-3669-1.
- Stanton, N.A. et al. (2006) 'Distributed situation awareness in dynamic systems: theoretical development and application of an ergonomics methodology', *Ergonomics*, 49(12–13), pp. 1288–1311. DOI 10.1080/00140130600612762.

- Star, S.L. & Ruhleder, K. (1996) 'Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces', *Information Systems Research*, 7(1), pp. 111–134. DOI 10.1287/isre.7.1.111.
- Staykova, K.S. & Damsgaard, J. (2016) 'Platform Constellations: The Case of KakaoTalk and LINE', *Proceedings of the Pacific Asia Conference on Information Systems (PACIS 2016)*, Paper 286.
- Talmar, M., Walrave, B., Podoyntsina, K.S., Holmström, J. & Romme, A.G.L. (2020) 'Mapping, analyzing and designing innovation ecosystems: The Ecosystem Pie Model', *Long Range Planning*, 53(4), 101850. DOI 10.1016/j.lrp.2018.09.002.
- Tavalaie, M.M. & Cennamo, C. (2021) 'In Search of Complementarities Within and Across Platform Ecosystems: Complementors' Relative Standing and Performance in Mobile Apps Ecosystems', *Long Range Planning*, 54(5), 101994. DOI 10.1016/j.lrp.2020.101994.
- Teece, D.J. (2007) 'Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance', *Strategic Management Journal*, 28(13), pp. 1319–1350. DOI 10.1002/smj.640.
- The Gazette (2016) 'The coffee house politicians: scandal and The London Gazette', 9 June. London: The Gazette.
- The London Gazette (1688/89) Advertisement referring to 'Mr. Edward Loyd's Coffee House in Tower Street', 18–21 February 1688/89.
- Tilson, D., Lyytinen, K. & Sørensen, C. (2010) 'Research Commentary—Digital Infrastructures: The Missing IS Research Agenda', *Information Systems Research*, 21(4), pp. 748–759. DOI 10.1287/isre.1100.0318.
- Tiwana, A., Konsynski, B. & Bush, A.A. (2010) 'Research Commentary—Platform Evolution: Coevolution of Platform Architecture, Governance, and Environmental Dynamics', *Information Systems Research*, 21(4), pp. 675–687. DOI 10.1287/isre.1100.0323.
- Toulouse School of Economics (2001) 'CSIO/IDEI, 2nd Joint Workshop on Industrial Organization', Toulouse, 1–2 June 2001. Conference programme.
- Toulouse School of Economics (2018) *The Origins of TSE*. Toulouse: Toulouse School of Economics.
- Tulane University (2011) 'Freeman School promotes two, appoints four to faculty for 2011–12', *Freeman News*, 30 August. New Orleans, LA: Tulane University.
- U.S. Department of Justice (2005) 'Statement regarding the closing of the Western Union Money Transfer Services investigation', Antitrust Division, 16 March. Washington, DC.

- U.S. Department of Justice (2024a) 'Michael L. Katz Biography'. Antitrust Division archive. Updated 2 January 2024.
- U.S. Department of Justice (2024b) 'Carl Shapiro Biography'. Antitrust Division archive. Updated 4 January 2024.
- UNESCO World Heritage Committee (2018) 'Decision 42 COM 8B.34: Göbekli Tepe (Turkey)'. 42nd session of the World Heritage Committee, Manama. Paris: UNESCO.
- University of Chicago (1949) 'Sears, Roebuck and Company', historical account published by the University of Chicago. Chicago.
- University of Chicago Library (2006) Guide to the Louis E. Asher Papers, 1894–1914. Hanna Holborn Gray Special Collections Research Center, University of Chicago Library.
- University of Chicago Library (2009) Guide to the Julius Rosenwald Papers, 1905–1963. Hanna Holborn Gray Special Collections Research Center, University of Chicago Library.
- Verhoeven, M. (2025) 'Relationality, immanence, hierarchy: The nature and culture of being(s) at Göbekli Tepe', *Cambridge Archaeological Journal*, 35(3), pp. 504–521. DOI 10.1017/S0959774325000113.
- Vogels, W. (2006) 'S3 – The Amazon Simple Storage Service', *All Things Distributed*, 14 March.
- Vogels, W. (2022) interview in 'The Evolution of Serverless at AWS with Dr. Werner Vogels', *Serverless Chats*, Episode 132, 11 April.
- Weiss, E., Wetterstrom, W., Nadel, D. & Bar-Yosef, O. (2004) 'The broad spectrum revisited: Evidence from plant remains', *Proceedings of the National Academy of Sciences*, 101(26), pp. 9551–9555. DOI 10.1073/pnas.0402362101.
- Western Union Company (2007) Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2006. Washington, DC: U.S. Securities and Exchange Commission.
- Western Union Telegraph Company (1868) *The Proposed Union of the Telegraph and Postal Systems: Statement of the Western Union Telegraph Company*. New York.
- Western Union Telegraph Company (c. 1820–1995) *Western Union Telegraph Company Records*, Collection NMAH.AC.0205. Archives Center, National Museum of American History, Smithsonian Institution, Washington, DC.
- Wolff, J.D. (2013) *Western Union and the Creation of the American Corporate Order, 1845–1893*. Cambridge: Cambridge University Press. DOI 10.1017/CBO9780511998041.

Wright, C. & Fayle, C.E. (1928) A History of Lloyd's from the Founding of Lloyd's Coffee House to the Present Day. London: Macmillan for the Corporation of Lloyd's.

Figure Credits and Rights

Figure 1.1 — Selected milestones and platform-firm market capitalisation. Author’s visualization. Market-capitalisation data are drawn from PwC’s *Global Top 100* series, based on Bloomberg data for 2009–2020 and S&P Global Market Intelligence for 2026. Platform-firm classification is date-specific and follows the transaction, innovation and hybrid-platform distinction in Cusumano, Gawer and Yoffie (2019). The research milestones beneath the quantitative series are selected orientation points in the intellectual history, not observations on the market-capitalisation axis and not a quantitative series of research activity.

Data record for Figure 1.1

Year	Classified platform firms appearing in the global Top 10	Combined market cap	Top-10 share
2009	Microsoft	0.163 tn USD	8.5%
2014	Apple; Google; Microsoft	1.196 tn USD	37.7%
2016	Apple; Alphabet; Microsoft; Facebook; Amazon	2.164 tn USD	58.5%
2019	Microsoft; Apple; Amazon; Alphabet; Facebook; Alibaba; Tencent	4.879 tn USD	80.2%
2020	Microsoft; Apple; Amazon; Alphabet; Alibaba; Facebook; Tencent	5.549 tn USD	69.9%
2026	Apple; Alphabet; Microsoft; Amazon; Meta	13.633 tn USD	57.0%

Figure 2.1 — Ohalo II and Göbekli Tepe. Ohalo II composite sickles: Groman-Yaroslavski, Weiss and Nadel (2016), via PLOS / Wikimedia Commons, CC BY 4.0. Göbekli Tepe main excavation area: German Archaeological Institute, photo E. Küçük, via Wikimedia Commons, CC BY 4.0.

Figure 3.1 — Edward Lloyd. Lloyd’s Coffee House, later depiction by William Holland (1789), via Wikimedia Commons; public-domain artwork / reproduction. *The London Gazette*, 21 February 1688; historical scan in the public domain.

Figure 4.1 — Western Union. Hiram Sibley (1807–1888), nineteenth-century engraved portrait, via Wikimedia Commons (file: *HiramSibley.jpg*); public-domain historical image. General Operating Department, Western Union Telegraph Building, New York, 1875, Library of Congress, Prints and Photographs Division; no known restrictions on publication.

Figure 5.1 — Richard Warren Sears. Richard Warren Sears (1863–1914), historical portrait, Wikimedia Commons, public domain. R. W. Sears watch advertisement, *Winchester Journal*, 5 January 1888, p. 3, public domain.

Figure 6.1 — Amazon Web Services. Werner Vogels at AWS Summit 2013, photograph by Raysonho, via Wikimedia Commons, CC0. Amazon S3 launch: screenshot of Jeff Barr,

“Amazon S3”, AWS News Blog, 14 March 2006; © Amazon Web Services. Reproduced here as documentary source material.

Figure 7.1 — Shopify. Tobias Lütke at Shopify+ in New York City, 2019, photograph by Benjamin Forrest / Good Thanks Media, via Wikimedia Commons, CC BY-SA 4.0. Snowdevil storefront, c. 2004, historical screen capture republished by Shopify, © Shopify. Copyright in the storefront image remains with its rights holder; no reuse permission for that image is asserted by this essay.

Figure 8.1 — Edward Lloyd across two Core Frames. Author’s analytical reconstruction based on the historical case and PBMC actor grammar. The wider maritime information and exchange frame is represented with distributed governance and does not imply one formal seventeenth-century Platform Owner.

Figure 8.2 — Western Union: same infrastructure, different transaction. Author’s analytical reconstruction based on the Western Union case. The figure distinguishes the telegraph-message and money-transfer CVUs while showing the continuity of network, offices, trust and accounting infrastructure.

Figure 8.3 — Sears-centred Platform Landscape. Author’s analytical reconstruction based on the Sears case. Sears is represented as Owner of the mail-order platform and as a user of the postal platform, while railway transport is shown separately as strategically consequential surrounding infrastructure rather than coded as one PBMC platform.

Figure 8.4 — Amazon internal services before AWS externalization. Author’s analytical reconstruction based on the AWS case. The figure represents recurring internal Provider–Consumer relationships around reusable technical capabilities before external market access becomes the defining feature.

Figure 8.5 — Shopify across commerce and agent-mediated commerce. Author’s analytical reconstruction based on the Shopify case. Question marks are intentional because Shopify’s role in the second frame depends on specification of the focal CVU and transaction.

Figure 9.1 — Adapted PBMC Core Frame. Author’s visualization. Adapted from the Platform Business Model Canvas actor grammar developed in Eisape (2019; 2020a; 2020b); the fixed geometry used for Landscape-level comparison is an adaptation introduced in this essay.

Figure 9.2 — Connected PBMC Core Frames. Author’s visualization. Conceptual representation proposed in this essay and informed by research on multiple-ecosystem participation and inter-platform ecosystems.

Figure 9.3 — Analytical recursion between PBMC and Platform Landscape. Author’s visualization. Conceptual representation proposed in this essay.