



Paper to be presented at the
35th DRUID Celebration Conference 2013, Barcelona, Spain, June 17-19

Architectural Control and Value Migration in Layered Ecosystems: The Case of Open-Source Cloud Management Platforms

Richard Tee

École Polytechnique Fédérale de Lausanne - EPFL

CDM

richard.tee@epfl.ch

C Jason Woodard

Singapore Management University

School of Information Systems

jwoodard@smu.edu.sg

Abstract

Our paper focuses on strategic decision making in layered business ecosystems, highlighting the role of cross-layer interactions in shaping choices about product design and platform governance. Based on evidence from the cloud computing ecosystem, we analyze how concerns about architectural control and expectations regarding future value migration influence the design of product interfaces and the degree of openness to external contributions. We draw on qualitative longitudinal data to trace the development of two open-source platforms for managing cloud-based computing resources. We focus in particular on the emergence of a layered "stack" in which these platforms must compete with both vertically integrated service providers and horizontally focused commercial software vendors. We find that the two platforms adopt distinctly different strategies, which presents an intriguing puzzle since their functionalities are nearly identical. Building on prior theory, we explain these differences in terms of the structure of interdependencies between each platform's lead sponsor and the firms operating in adjacent layers of the ecosystem. While our findings are preliminary due to the ongoing nature of the research, we speculate about their implications for the literatures on modularity and innovation, technology and industry evolution, and strategy in business ecosystems.

Architectural Control and Value Migration in Layered Ecosystems: The Case of Open-Source Cloud Management Platforms

Richard Tee
Postdoctoral Research Fellow
College of Management of Technology
École Polytechnique Fédérale de Lausanne (EPFL)
richard.tee@epfl.ch

C Jason Woodard
Assistant Professor
School of Information Systems
Singapore Management University
jwoodard@smu.edu.sg

Abstract¹

Our paper focuses on strategic decision making in layered business ecosystems, highlighting the role of cross-layer interactions in shaping choices about product design and platform governance. Based on evidence from the cloud computing ecosystem, we analyze how concerns about architectural control and expectations regarding future value migration influence the design of product interfaces and the degree of openness to external contributions. We draw on qualitative longitudinal data to trace the development of two open-source platforms for managing cloud-based computing resources. We focus in particular on the emergence of a layered "stack" in which these platforms must compete with both vertically integrated service providers and horizontally focused commercial software vendors. We find that the two platforms adopt distinctly different strategies, which presents an intriguing puzzle since their functionalities are nearly identical. Building on prior theory, we explain these differences in terms of the structure of interdependencies between each platform's lead sponsor and the firms operating in adjacent layers of the ecosystem. While our findings are preliminary due to the ongoing nature of the research, we speculate about their implications for the literatures on modularity and innovation, technology and industry evolution, and strategy in business ecosystems.

¹ We would like to thank interviewees for taking the time to share their views; Abhishek Kumar for excellent research assistance; and participants at the OUI 2012 at MIT/HBS and SMS Lake Geneva Conference 2013 for feedback. All errors are ours.

1. Introduction

In industries that produce modular systems, architectural choices have strategic consequences (Langlois and Robertson, 1992; Baldwin and Clark, 2000; Schilling, 2000). These consequences are especially acute for system components that serve as platforms for complementary products and services (Bresnahan and Greenstein, 1999; Gawer and Cusumano, 2002; Iansiti and Levien, 2004). Architectural control is a particularly salient issue in platform-based competition (Morris and Ferguson, 1993), where the ability to create and capture economic value depends heavily on striking an appropriate balance between “open” and proprietary strategies (West, 2003; Boudreau, 2010; Parker and Van Alstyne, 2013).

Consider the decision by Twitter, the operator of the popular social networking service, to provide an application programming interface (API) to third-party software developers.² This move, which came only months after the service launched in 2006, allowed developers to extend Twitter’s functionality, which in turn increased the attractiveness of the service to the company’s rapidly growing user base. Dozens of new software clients and Twitter-enabled web applications soon appeared, complementing Twitter’s official offerings and opening the Twitter platform to broad participation by a variety of firms.³ However, as the company’s strategy shifted toward generating revenue from advertisements, the third-party clients began to pose a threat as they competed with Twitter’s official clients but did not allow Twitter to directly control the display of ads. In response, Twitter revised its API and terms of service in mid-2012 to impose both technical and legal barriers to high-volume clients.⁴ This example illustrates the delicate balance

² <http://blog.twitter.com/2006/introducing-twitter-api>

³ <http://blog.twitter.com/2010/twitter-platform>

⁴ <http://dev.twitter.com/blog/changes-coming-to-twitter-api>

firms must strike between openness and control, and the strategic impact of seemingly minor design decisions.

While existing literature has made progress in understanding the structure and dynamics of platform competition, much of the prior theoretical work focuses on the relationships between a single platform and its complementors (e.g., Parker and Van Alstyne, 2005) or between two competing platforms (e.g., Rochet and Tirole, 2003; Casadesus-Masanell and Yoffie, 2007; Zhu and Iansiti, 2011). These are no doubt important cases, both in theory and in practice, but many settings are characterized by multiple platforms (and potential platforms) that interact with each other through loosely coupled interdependencies in a layered modular architecture (Gawer and Henderson, 2007; Yoo et al. 2010). In these settings, some of the strategic levers emphasized in the theoretical literature, such as pricing and cross-subsidization, are of secondary importance to the basic concern of staking out a viable position in a crowded and confusing competitive space. Moreover, the standard assumption of rational profit-maximizing behavior is simply untenable due to the high degree of uncertainty and limited foresight of the players.

These challenging environments are well described by the literature on nascent markets and sensemaking (Kaplan and Tripsas, 2008; Santos and Eisenhardt, 2009), as well as the concept of an “era of ferment” from the innovation literature (Anderson and Tushman, 1990). Likewise, the dynamic capabilities literature has highlighted the importance of learning and adaptation when competing in complex environments where high interdependence creates uncertainty about the strategic impact of a firm’s actions (Teece et al., 1997; Eisenhardt and Martin, 2000; Helfat and Peteraf, 2003). The growing literature on business ecosystems has also articulated principles for

managing relationships between stakeholders, ensuring the “health” of the ecosystem, and matching a firm’s strategy to its ecosystem (Moore, 1996; Iansiti and Levien, 2004; Adner, 2006). While all of these literatures contribute valuable insights for firms in multi-platform, multi-layered ecosystems such as mobile communication, digital content distribution, social networking, and enterprise application software, they do not address the specific strategic challenges posed by competing in an evolving modular architecture. Despite the significant attention devoted to these ecosystems by academic researchers, industry practitioners, and the popular press, we still lack robust insights about how their layered structure emerges or what firms can do to influence their evolution in pursuit of sustainable competitive advantage.

To address this gap, this paper presents an inductive study of firms’ strategic decision making in the cloud computing ecosystem. The layered architecture of this ecosystem has begun to stabilize at a high level of abstraction (Liu et al., 2011) but remains very much in flux at the more concrete level that is strategically relevant to firms engaged in developing cloud-based products and services. Drawing on evidence from an ongoing case study of two open-source software platforms for managing cloud computing resources, we analyze how the two platforms’ corporate sponsors strategically shape the design of their architectural neighborhood: the layers in which they participate directly, the adjacent layers linked by technological interdependencies, and the layers in which their competitors and complementors participate. We use qualitative longitudinal data to trace the emergence of a layered “stack” in which the focal platforms must compete with both vertically integrated service providers and horizontally focused commercial software vendors.

We find that the two platforms adopt distinctly different strategies, which presents an intriguing puzzle since their functionalities are nearly identical. In the remainder of the paper, we seek to explain these differences in terms of observable characteristics of the firms and software artifacts involved. Consistent with prior theory and a sizable body of practitioner literature, we find that architectural control—especially the threat of control by others—is a salient concern for both platform sponsors, although the focus of their concern differs due to the fact that they participate in different layers of the cloud architecture. While we find little evidence of strategic pricing behavior, we observe that expectations about future value migration (Slywotzky, 1996; Jacobides and Tae, 2012) play an important role in the firms’ strategic decisions—in particular, the extent to which certain layers may become commoditized or “de-commoditized” over time (Christensen and Raynor, 2003). Finally, we examine the strategic levers by which the firms attempt to shape their ecosystem to their advantage. As in Garud et al.’s (2002) study of Sun Microsystems’ Java technology, institutional entrepreneurship plays a central role in both firms’ strategies. We also observed “architectural entrepreneurship” in the form of design moves intended primarily to influence the technological trajectory of the ecosystem rather than achieve a functional goal.⁵

Although these findings are still preliminary due to the ongoing nature of the research, we believe they can contribute to the existing literature by shedding light on the antecedents of strategically significant architectural design choices. While further research is needed to understand the consequences of these choices at the firm level, their impact on the cloud computing ecosystem is already becoming apparent. This suggests that it may be fruitful to

⁵ Note that our conceptualization of architectural entrepreneurship differs from Richard and Devinney (2005), who focus on “the firm’s architectural knowledge to cast the supply chain more efficiently than would be achieved through the price-driven market”.

theorize more generally about the role of strategic design moves in the emergence of layered modular architectures.

The rest of the paper is organized as follows. Section 2 describes our data sources and research methods. Section 3 describes our research setting in more detail, providing background on the two focal platforms we studied (OpenStack and CloudStack) and brief profiles of the firms that play key roles in the architectural neighborhood of these platforms (Amazon.com, VMware, Rackspace, Citrix, and Eucalyptus). Section 4 presents our case evidence, including a summary of the institutional and architectural design moves we observed, and interpretive comments from our industry informants. Section 5 concludes by summarizing our findings and discussing their implications, as well as the limitations of the study and our plans for future work.

2. Data Sources, Setting, Methods

Our study focuses on software development in an emerging technological field known as cloud computing. Generally, cloud computing refers to computing resources being delivered as a network-based service. We chose this as our overall setting for several reasons. First, services in the area of cloud computing can be represented as a layered modular architecture (or “stack” in software engineering terms), which allows us to systematically observe how design choices affect the overall technological architecture. Second, although the cloud computing ecosystem is evolving rapidly, much of this evolution occurs through observable design moves that result in discrete changes to the stack architecture that can be traced over time. The process is more tractable in open source software initiatives (our main focus), as the overall development process is documented and readily accessible. Third, the dynamics of value migration in this ecosystem

are at a relatively early stage, providing a useful contrast to more mature industry ecosystems such as personal computers, automobiles, and semiconductors.

Our case evidence is drawn from a combination of archival and interview data. We collected several types of archival data: press releases of relevant firms, official company weblogs, as well as a selected set of industry-focused websites covering cloud computing.⁶ Using an approach similar to Woodard et al. (2013) we coded events to create a systematic overview of the various design moves made by the firms we observe. This analysis was an iterative process consisting of several steps. First, one of the authors collected a set of articles related to “infrastructure-as-a-service” cloud computing initiatives. Next, this set of articles was analyzed by a research assistant, focusing on several dimensions: a summary of the event, the firms involved in this event, its strategic rationale underlying, and implications for technical compatibility. After this information was collected in a spreadsheet, the authors then selected several key events to analyze in further detail, providing the main input for the empirical section of the paper. The set of selected events co-evolved with our general theoretical framework and our understanding of the setting. The initial focus of the project centered broadly on ecosystem strategies in open-source cloud computing initiatives. As the project continued, we narrowed the focus to the issues of standardization and compatibility, and value migration.

In addition to our archival data, we have conducted several interviews to date with informants actively involved in the design of and strategy concerning the two key initiatives we focused on.

Interviews were done face-to-face when possible, and through video conferencing or telephone

⁶ The industry-focused websites included CloudAve, Forbes CIO Network, Forrester Research, Gartner Blog Network, GigaOM Cloud, ReadWriteCloud, and Talkin’ Cloud. By using a variety of sources, we aimed to ensure coverage of the major events relevant to our setting.

when a personal meeting was not possible. The interviews, which were transcribed by a professional transcription firm, served two main purposes. First, they allowed us to get a better understanding of the motivations underlying the decisions we observed through public sources. Individual quotes (anonymized as agreed with interviewees) were also included to illustrate firm's decision making. Second, it enabled us to discuss our understanding of various events with informants and adjust our overall framework where necessary. The interviews, combined with our archival data, also helped us in building the overall layered stack depicted in Figure 1.

3. Empirical Setting: Cloud Management Platforms

Our empirical setting examines an emerging market in the computer industry referred to as "cloud computing". One well-known example of cloud computing is Dropbox, which provides firms or individual end-users with network based storage. Our study focuses on a less visible segment of cloud computing referred to as "infrastructure-as-a-service" (typically abbreviated as "IaaS"). Within IaaS we focus on two initiatives offering IaaS as an on open-source development project, OpenStack and CloudStack, led by two different firms (Rackspace and Citrix respectively).

We focus on in particular on the software layers responsible for orchestrating cloud-based computing resources such as processing power, network bandwidth, and storage capacity. This software layer is also referred to as Cloud Management Platforms (CMP).⁷ At the time of

⁷ Gartner, a well known IT research company defines Cloud Management Platforms as "integrated products that provide for the management of public, private and hybrid cloud environments. The minimum requirements to be included in this category are products that incorporate self-service interfaces, provision system images, enable metering and billing, and provide for some degree of workload optimization through established policies. More-advanced offerings may also integrate with external enterprise management systems, include service catalogs, support the configuration of storage and network resources, allow for enhanced resource management via service

writing, several commercial products and commercially sponsored open-source projects were competing for adoption. The two most prominent open-source projects, OpenStack and CloudStack, are sponsored by firms with very different commercial interests. OpenStack is sponsored by Rackspace, a hosted services provider that sees Amazon.com as a primary competitor. CloudStack is sponsored by Citrix, a software company that perceives the VMware division of EMC as a greater threat.

Particular to both initiatives is the goal to define the entire “value stack”, in contrast to most other industry consortia that focus on a particular subsystem (or layer). Therefore, our research setting is distinctive in two ways. First, the initiative is not confined to a single product category or standard but spans multiple architectural layers. Second, the interactions within the ecosystem take place against a backdrop of significant competitive pressures, both within layers and with respect to vertically integrated rivals.

To better understand our two focal projects, we provide a brief sketch of the competitive landscape in which CloudStack and OpenStack are situated. Three offerings, Amazon Web Services (also referred to as AWS) and VMware (in particular their vCloud product suite), are especially relevant to our focal projects. While Amazon is known mostly for its online retail activities, it has also built a considerable business in providing web services to other firms. Its AWS portfolio comprises a variety of products; the two most relevant ones in this context are Elastic Compute Cloud (EC2) and Simple Storage Service (S3). Amazon EC2 lets their customers run software applications without having to set up their own physical computers.

governors and provide advanced monitoring for improved “guest” performance and availability.”
<http://www.gartner.com/it-glossary/cloud-management-platforms> (May 30, 2013)

Instead, EC2 delivers additional virtual servers as required, and charges fees based on time and resource usage. Amazon S3 lets customers store data remotely, again sidestepping the need for physical storage. For both services, the customers are usually not individual end-users but other businesses, such as Dropbox and Netflix.

Besides Amazon, VMware (owned by EMC) is another relevant competitor in cloud computing. Generally, VMware focuses on virtualization software. This allows customers to run, for example, a Windows operating system on an Apple computer. Similar products are also available for server and storage hardware, facilitating storage in local corporate data centers known as “private clouds” (as the data is stored on premise, not in a remote (i.e. “public”) location). VMware offers a proprietary cloud computing operating system called vCloud. It joined OpenStack in October 2012, and extended OpenStack support for its ESX hypervisor (which appears to be done in part to drive sales for VMware’s other products, such as vSphere).⁸

Figure 1 depicts the IaaS segment of the cloud computing ecosystem. The major layers, from bottom to top, include hosting services, hardware, virtualization software (also called hypervisors), resource orchestration, and the application programming interfaces (APIs) used by developers of cloud-aware applications and services.

Finally, another CMP is Eucalyptus, offered by Eucalyptus Systems Inc. Eucalyptus started as a university research project at the University of California, Santa Barbara in 2007, and was commercialized as an open source company in 2009. The Eucalyptus cloud platform is compatible with various Amazon AWS services, including EC2 and S3.⁹

⁸ <http://www.networkworld.com/news/2012/101712-vmware-openstack-263473.html>

⁹ <http://www.eucalyptus.com/faq> and <http://www.eucalyptus.com/about/story>

Main cases: OpenStack and CloudStack

Our focal cases are OpenStack and CloudStack, two open source CMPs sponsored by two commercial firms (Rackspace and Citrix respectively). In 2010, Rackspace (jointly with NASA) launched an initiative called OpenStack to develop a suite of open-source software for the resource orchestration layer, which provides tools to manage cloud-based computing resources such as virtual machines, object storage, and network connectivity. Initially unbeknownst to each organization, both NASA and Rackspace were developing technologies similar to Amazon's EC2 and S3. NASA was working on a project called Nova, which basically mimicked Amazon's EC2 service. In turn, Rackspace's Swift project was similar to Amazon's S3. Coincidentally, project members from each team connected, and subsequently decided to work together as part of the OpenStack project.¹⁰

Our second focal project, CloudStack, began as a software suite developed by a startup company called Cloud.com, which was purchased by Citrix in 2011 and donated to the Apache Foundation in 2012. Citrix also acquired the Xen hypervisor (which it subsequently released in open-source and commercial versions) with its purchase of XenSource in 2007.

<< insert Figure 1 about here >>

Figure 1 provides a schematic overview of the IaaS ecosystem, focusing on the firms' services summarized above. In this figure arrows indicate dependence; nesting indicates compatibility (i.e., nested boxes are interchangeable with respect to modules that depend on their enclosing

¹⁰ See <http://www.wired.com/wiredenterprise/2012/04/openstack/> for more details on OpenStack's founding.

boxes). Qualitatively, the figure shows that Amazon Web Services is a fully integrated set of offerings that include hosting services, while VMware offers separate products at each layer except hosting; these products may be purchased separately or bundled.

4. Key Strategic Moves in Architecture and Governance

This section focuses on the key strategic moves Rackspace and Citrix made in relation to their respective CMPs. Figure 2 provides a timeline of key events related to our focal projects (CloudStack and OpenStack) and the competitive landscape in which they operate.

<< insert Figure 2 about here >>

OpenStack launch

OpenStack generated widespread publicity in the industry upon its announcement. One interviewee previously working with OpenStack reflects on its momentum as follows:

I suspect it had something to do with the momentum that OpenStack had, but OpenStack is this crazy, crazy phenomenon that has the most buzz generated about it for anything anyone's ever seen. *So it's really kind of shocking, and surprising, and exciting all at the same time.*

Motivation underlying Openstack launch

An obvious question that arises is, why would Rackspace, a pure-play hosting services company, invest in an open-source cloud orchestration project? The stack represented in Figure 1 suggests

two plausible motives: (a) to draw customers away from its direct competitor Amazon's fully integrated offerings, and (b) to make it more difficult for any other firm (notably Citrix or VMware) to secure a dominant position in the orchestration layer.

The first motive is consistent with the idea that open standards can soften competition between vertical bundles (cf. Baldwin and Woodard, 2007); by making it possible for other firms to mix and match orchestration software with modules from other layers, Rackspace avoids the direct escalation of price competition that would have occurred if it had simply launched its own proprietary alternative to Amazon Web Services.

The second motive, known in the software industry as “commoditize your complements” (cf. Shapiro and Varian, 1998, p. 279), raises more difficult questions. This tactic is equivalent to eliminating a layer of the architecture because prices in that layer get driven toward zero, which diminishes their ability to affect pricing decisions in other layers. System prices fall, but aggregate industry profits may go up or down depending on how the pre- and post-commoditization system prices compare to the price that would be set by a fully integrated monopoly (“one big firm”). Moreover, when the intent of commoditization is to deprive a firm's complementors of profits that could be used to attack it (e.g., through “platform envelopment”, cf. Eisenmann et al., 2011), the result tends to promote the emergence of layer monopolies.

One interviewee reflects on potential commoditization of lower levels in the stack, such as hypervisors, as follows:

Amazon came and showed everyone how you can make money off the cloud and use the cloud, people came around and started selling their own mostly open core, open source *model type cloud solutions (...). And we're getting to where there's actually quite a few orchestration layers out there for people to choose from and they're all open source. They all work with all the hypervisors. They're all [...] becoming a commodity.*

Awareness of commoditization of existing components in the stack may prompt a strategic reorientation for some firms. On the issue of value migration, another interviewee discusses this in terms of the perceived longer term value of service delivery:

So I think the ultimate value of this is in the customer experience that's delivered at the top, and, of course, then there's the operational efficiencies and so on that one can get through the horizontal scalability, and the automation provided through the platform, but I think that part will be commoditized faster. And the part that will remain the most valuable is, what are the organizations and technology combinations, together, that can deliver the best service for the customers? I mean, that's where the real value's going to end up being.

Motivation underlying CloudStack launch

To complicate matters further, Citrix and Rackspace appear to be engaging in similar strategies but with important differences. Rackspace has made no secret of its view that Amazon represents the biggest threat to open standards for cloud computing, which led it to develop its own API for

OpenStack.¹¹ Citrix, on the other hand, is perceived by industry observers to fear VMware more than Amazon, leading it to adopt Amazon's APIs but continue to invest in a commercial hypervisor (Xen) to rival VMware's (ESX) — even though the major hypervisors are all largely compatible with each other, leaving little room for architectural lock-in.¹² These differences, which led Citrix to drop its support for OpenStack when it contributed CloudStack to Apache,¹³ cast doubt on the idea that any simple coordination mechanism (e.g., a traditional standards-setting organization) could bring about convergence to the kind of profitable open-standards equilibrium envisioned in the analytical model introduced in Baldwin and Woodard (2007). And yet, these are hardly scorched-earth tactics either. It thus remains plausible that the overall cloud computing ecosystem will ultimately avoid the extremes of vicious price wars (ruinous in the short term) or “Wintelism” (Borras and Zysman 1998). The latter is potentially ruinous in the longer term as responsibility for innovation becomes too concentrated, creating openings for new entrants like Google and vertically integrated rivals like Apple.

Competitive dynamics and expectations

Competition between the various cloud offerings is not straightforward, given differences not only between proprietary and open-source based solutions, but also differences in the activities of the stack the various firms are engaged in. One interviewee expects that in relation to open source based cloud solutions, one initiative will prevail:

¹¹ http://blogs.gartner.com/lydia_leong/2012/04/06/ecosystems-in-conflict-amazon-vs-vmware-and-openstack/#comment-7859

¹² http://blogs.gartner.com/lydia_leong/2012/04/06/ecosystems-in-conflict-amazon-vs-vmware-and-openstack/#comment-7859

¹³ <http://gigaom.com/cloud/5-takeaways-from-the-cloudstack-openstack-dustup/>

We think there will be probably a Microsoft-driven cloud, probably another proprietary cloud, possibly VMware, and we think there will be one open-source cloud standard. And *we want to be that latter one (...). So we think there's definitely room for a couple of proprietary offerings, and one open-source offering. We think one. We think probably not two, because we can't see a reason why the market would support two competing open-source cloud offerings, rather than having the advantage of having one standardized open-source offering.*

Comparing CMPs: governance differences

Besides these cross-layer interdependencies, we observe important differences between CloudStack and OpenStack in terms of organization (e.g., project governance) and technology (e.g., APIs). Citrix turned over the governance of its project to the Apache Foundation, a well-established forum for independently managed open source projects. Furthermore, it chose to implement Amazon's APIs, in particular in relation to two key components of Amazon's Web Services, EC2 and S3. This has diminished technological independence of CloudStack, and by extension Citrix, as it does not directly control the development of interfaces for future development of the platform. By contrast, OpenStack's governance was initially mostly led by Rackspace. Only recently it has introduced an independent—albeit tiered—foundation to manage the overall project, allowing other firms to influence the direction of the project.

Comparing CMPs: interface (API) differences

In terms of design choices, the difference in implementation between OpenStack and CloudStack and their compatibility with Amazon's APIs reveals several important differences. Given the

widespread diffusion of AWS, it is important for other services, including CloudStack and OpenStack, to be compatible with Amazon's APIs. However, a close inspection of how these projects have realized compatibility with AWS, shows subtle differences with important implications. While both services are compatible, OpenStack's APIs are not "natively" compatible. Instead, OpenStack's native APIs are designed in a way that it "offers non-native 'AWS API Compatibility Modules' for EC2 and S3".¹⁴ As a result of this design choice, OpenStack is more autonomous in terms of future development in case Amazon makes technical changes to its APIs, or if it changes the conditions surrounding their usage. On the differences between CloudStack and OpenStack, one informant perceives the issue of API compatibility as follows:

With CloudStack now being open, I think there's a number of major differentiators; one is that it's aligned in terms of API compatibility with AWS, which is the world's biggest cloud and OpenStack isn't—that's a key differentiator.

Another interviewee further illustrates how firms engage with Amazon's APIs as follows:

Amazon is the de facto standard, and many organizations, technologies try to provide a translation layer that allows the Amazon syntax to be used across to their own API so CloudStack does that. There's a number of API abstraction projects out there that allow you to have a common API in translation but create, destroy, reboot instances across cloud providers is important.

¹⁴ http://blogs.gartner.com/lydia_leong/2012/04/06/ecosystems-in-conflict-amazon-vs-vmware-and-openstack/#comment-7859

Compatibility with AWS (or “Amazon-style architectures”) has been used to promote cloud services. However, how exactly compatibility is defined and how it applies to the various cloud services is an area of contention.¹⁵ Besides CloudStack and OpenStack, other services have also emphasized their compatibility with AWS. For example, in March 2012 Eucalyptus announced an agreement with Amazon that:

... enables customers to more efficiently migrate workloads between their existing data centers and AWS while using the same management tools and skills across both environments. As part of this agreement, AWS will support Eucalyptus as they continue to extend compatibility with AWS APIs and customer use cases. Customers can run applications in their existing datacenters that are compatible with popular Amazon Web Services such as Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3).¹⁶

The strategic importance of APIs was further emphasized by another interviewee:

But building a real ecosystem around the APIs, so that it's very difficult for anyone else to compete, I think that's a sustainable advantage. And secondly, building the combination of technologies and the organization that can deliver the best customer

¹⁵ http://blogs.gartner.com/lydia_leong/2012/04/06/ecosystems-in-conflict-amazon-vs-vmware-and-openstack/, <http://gigaom.com/2012/04/07/true-or-false-citrix-is-more-compatible-with-aws/>

¹⁶ <http://www.eucalyptus.com/news/amazon-web-services-and-eucalyptus-partner>, <http://www.forbes.com/sites/danwoods/2012/04/04/questions-amazon-should-answer-about-its-cloud-strategy/>

service, I think that's another sustainable competitive advantage that you'll get. I mean, that's kind of in the longer term.

Antecedents driving interface (API) differences

We can distinguish several antecedents to these decisions. Some can be traced back to technological differences; more importantly, other differences seem to be rooted in the firm's position in the innovation ecosystem and their ability to capture value. From a technological perspective, there are important differences in the histories of the projects. CloudStack is based on a commercial product (initially known as Cloud.com) built by a small entrepreneurial firm. As such, it has a more monolithic source code compared to large scale distributed projects (cf. MacCormack et al. 2012). By contrast, from the outset OpenStack has been based on subsystems developed by a variety of organizations (in particular Rackspace and NASA). As such, the technological architecture has been more modular from the beginning. While these technological differences can explain some of the contrasts highlighted above (e.g. project governance), strategic motivations appear to be an overriding factor. In particular, as mentioned earlier, there are important differences in the location in the ecosystem of the firms backing CloudStack and OpenStack. As a result, they compete with different firms in the ecosystem and also use different ways to create and capture value. OpenStack's main sponsor, Rackspace, competes primarily with Amazon and their various web services. By contrast, CloudStack's main sponsor, Citrix, competes primarily with VMware's products.

Overall, our data reveal a dynamic, highly interdependent setting where participants show divergence both in their perception of the competitive landscape and their resulting strategies. In

turn, these strategies are reflected in the design choices firms make, affecting both the structure of the business ecosystem as well as the longer term evolution of the industry. Table 1 summarizes the design moves made by each project in terms of the timing of various key events and their underlying motivations.

<< insert Table 1 about here >>

5. Discussion and Concluding Remarks

At least since Teece's (1986) work on complementary assets and their role in how firms profit from innovations, we know of the strategic importance of interdependence. More recent work has extended Teece's framework by highlighting how firms might strategically shape the structure of complementarity, for example by becoming an industry "bottleneck", thereby facilitating value capture (Jacobides et al., 2006). Empirical work in this area has also shown how the location of technological bottlenecks might affect innovation, with important implications for market share of incumbents and challengers (Adner and Kapoor, 2010). Our study extends existing work by considering how firms strategize in dynamic, complex environments, such as the multi-layer stack that characterizes our empirical setting. In particular, we have focused on the drivers of a firm's architectural design choices. Our case suggests that in such interdependent settings, firms' design choices are far less straightforward compared to settings where the competitive parameters are fewer and more clearly defined.

In particular, we highlight two related issues that arose from our inductive study on the drivers of a firm's architectural design choices. First, we observe how differing expectations regarding

industry-wide value migration is reflected in variations in design moves. In multi-layered stacks, different firms are (by definition) located in distinct parts of the value chain. As such, assuming a firm's strategy is path-dependent (Nelson and Winter, 1982), we also observe how their subsequent design moves differ according to their existing position in the industry architecture. Second, and in relation to the issue of value migration, such expectations in turn inform choices regarding compatibility with existing services. As our empirical analysis shows, firms have a range of options in relation to compatibility. Most interestingly in this setting, besides the choice of being at all compatible with services from other firms (some of which are direct rivals, while others serve more as complements), firms can also choose how to be compatible with other services. While the economic impact of compatibility is well known (see e.g. Farrell and Saloner, 1985; Katz and Shapiro, 1985), there has been much less work on the strategic design of standards and compatibility choices, individual exceptions notwithstanding (e.g. Besen and Farrell, 1994; Garud et al., 2002; Woodard and West, 2011). Our study points to the fact that choosing how to be compatible may be a strategic choice itself. A related perspective on architectural design has also been introduced in recent work that has pointed to decomposition as a strategic variable (e.g., Baldwin and Henkel's (2011) work on intellectual property and modularity). While such differences may seem relatively minor, they reflect differences in firms' perception of the evolution of the competitive landscape, and can have important implications for lock-in, innovation and value capture.

Overall, the firm behaviors we observe appear to be driven by two related factors, architectural control and expected value migration. First, firms wish to retain architectural control by ensuring that adjacent layers in the stack are not controlled by a single dominant vendor. As the firms

sponsoring the focal cases are located in different parts of the ecosystem, their competitive responses differ in key technological choices (e.g. interface compatibility). Second, as firms observe how value migrates to different parts of the stack, they attempt to commoditize those layers on which their ability to capture value depends. The various design moves we observe can be characterized as “architectural entrepreneurship”: actions intended primarily to strategically influence the technological trajectory of the ecosystem rather than achieve a particular functional goal. While our findings are preliminary due to the ongoing nature of the research, we offer several speculations about their implications for the literatures on strategy in business ecosystems, technology and industry evolution, and modularity and innovation.

First, our study may provide several implications for the emerging literature on business ecosystems and value migration. The idea that designers “see and seek value” is axiomatic in the literature on design evolution and modular systems (cf. Baldwin and Clark, 2000). However, few studies have, to our knowledge, analyzed how firms strategize to attain value in complex business ecosystems such as our setting. By analyzing in-depth decision making, our study highlights how value might migrate across different layers due to commoditization and de-commoditization; how firms compete through strategic design moves, in particular in relation to compatibility; and how standards might be beneficial (in terms of value creation and value capture) for firms operating in such innovation ecosystems (Adner and Kapoor, 2010).

Second, our findings might also inform the literature on industry evolution and technology lifecycles. Our setting is characterized by endogenous innovation (i.e., not arising out of processes exogenous to an individual firm, such as technological breakthroughs or new

knowledge¹⁷). Instead of a “natural” co-evolution characterized by selection of competing variants (Anderson and Tushman, 1990; Murmann and Frenken, 2006), we see pro-active attempts to shape the overall design of the architecture, instead of more straightforward competition between architecturally similar systems. The dynamics observed in our setting also clearly depart from the model depicted in the disruptive innovation literature (Christensen, 1997), which contrasts the strategies of resource rich incumbents to resource constrained entrepreneurial firms. In contrast, no firm in our setting is clearly an incumbent or challenger, as each is active and leading in different parts of the stack.¹⁸ As such, the competing initiatives we observe involve more complex behavior based on judicious strategies in relation to technological compatibility and organization. For example, some of the initiatives we observed attempt to shape the industry trajectory by promoting technology that avoids lock-in through proprietary standards. Such strategies may in part reflect learning on behalf of the various industry participants, some which have suffered from existing industry architectures (Jacobides et al., 2006), for example the dominance of Microsoft in the PC industry.

Third, our paper may contribute to recent work that has studied the alignment between the design of products and organizations. In particular, recent work has invoked the notion of “mirroring” between a product’s architecture and its subsequent organization (Henderson and Clark, 1990; Colfer and Baldwin, 2010; Cabigiosu and Camuffo, 2012). This case highlights how strategy might be reflected in organizational (especially in relation to project governance) and product

¹⁷ In contrast, the technologies and services observed in our setting arise from incremental improvements (e.g., in network speed and availability), not a clear technological discontinuity.

¹⁸ To give an indication of the respective sizes of our focal cases, the total revenues in 2012 were: Citrix \$2.20 bn (2011), Rackspace \$1.31 bn, VMware \$4.61 bn (its parent company EMC \$20.01 bn), and Amazon \$61.09 bn (however its cloud business has been valued at less than \$10 bn, see e.g. <http://seekingalpha.com/article/1038151-what-is-amazon-web-services-really-worth>).

choices (in particular technological dependence in terms of compatibility). However, the existing technological architecture does not necessarily mirror governance choice. In our case the more modular project (OpenStack) was initially characterized by a governance model dominated by a single firm, whereas the more integrated architecture (CloudStack) had a more distributed governance model. These results add further support for the idea that mirroring is contingent, and cannot simply be derived from technological structure. While in some settings knowledge has been suggested to be a key driver of product and organizational architecture (Brusoni and Prencipe, 2006), our setting highlights the role of compatibility and value migration.

Finally, it is important to note that our study has several important limitations. First, as an inductive, qualitative study our results may not be generalizable to other settings. Depending on data availability, future work might quantitatively test our findings in other settings. Second, we have purposively chosen a fast-moving setting. As a result, we have little conclusively knowledge in terms of competitive outcomes. Therefore, it is important to note that our findings have no normative implications for firm strategy, for example in relation the way in which compatibility is achieved. Future work may also examine how the issue of compatibility plays out in settings beyond software development.

Figures and Tables

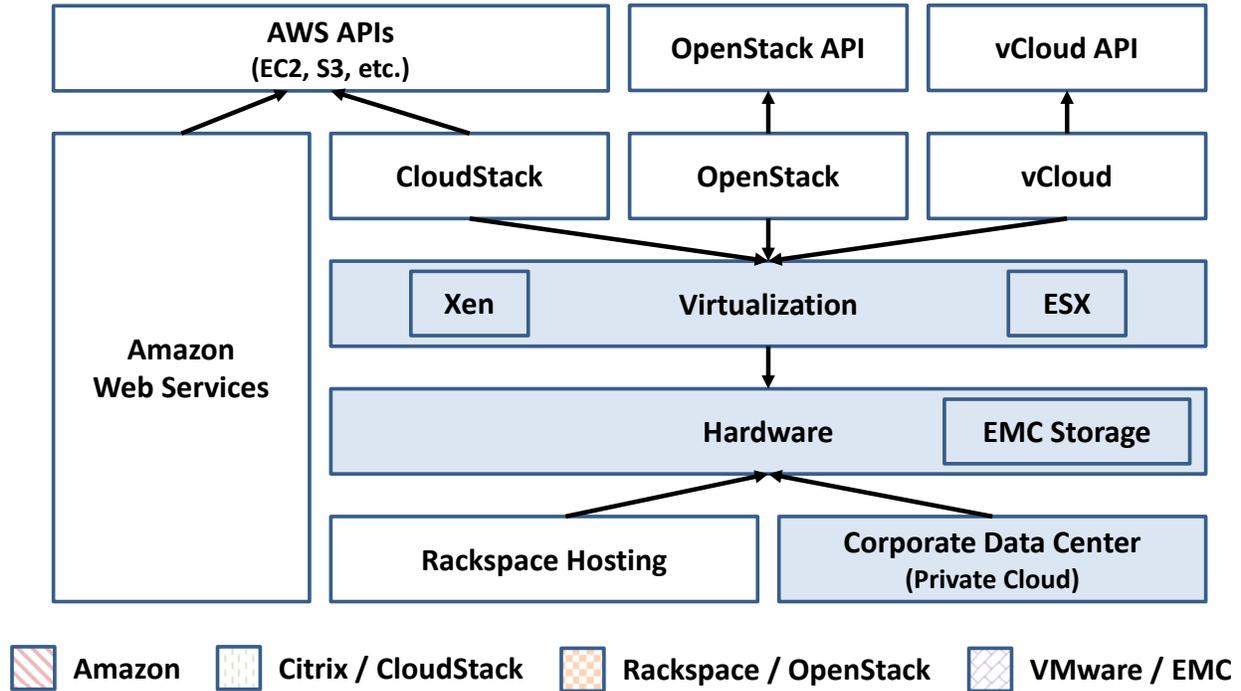


Figure 1: Layered structure of the infrastructure-as-a-service (IaaS) segment of the cloud computing ecosystem

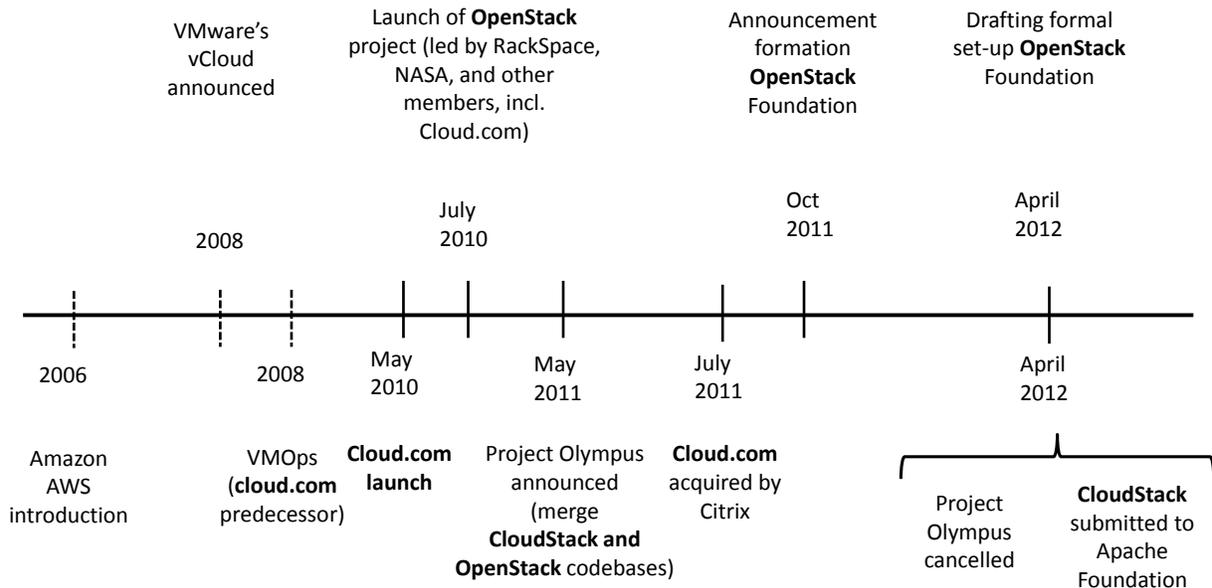


Figure 2: Timeline of key events related to CloudStack and OpenStack

OpenStack – Design Moves			
Date	Event	Description	Motivation
July 2010	Creation of OpenStack (Co-founded by Rackspace and NASA)	Facilitate development of open-source cloud orchestration software	Allows variety of stakeholders to compete with Amazon’s increasing dominance of cloud computing market
	Compatibility with AWS API’s	CloudStack’s APIs independent of Amazon’s (non-native compatibility)	Decrease dependence on third-party API’s (including AWS)
May 2011	Project Olympus initiated	Merging the codebases of the Cloud.com and OpenStack projects	Combine strengths of respective projects and build market share
April 2012	Cancellation Project Olympus	Merging of codebases halted due to CloudStack submission to Apache Foundation	Citrix focus on CloudStack as more mature solution
April 2012	OpenStack foundation initiated	Creation of OpenStack Foundation with tiered membership levels	Decrease (perceived) dominance of Rackspace governance
Oct 2012	VMware joining OpenStack	VMware joins OpenStack as gold member	Joining of cross-layer competitor
Oct 2012	ESX hypervisor support	VMware extends support for ESX hypervisor in OpenStack	VMware wants OpenStack customers to use vSphere
CloudStack – Design Moves			
Date	Event	Description	Motivation
May 2010	Cloud.com launch	Development of commercial cloud orchestration software	Commercial rival to Amazon and other cloud computing solutions
July 2011	Cloud.com acquisition by Citrix	Cloud.com team and product portfolio acquired by Citrix	Allows Citrix to enter cloud orchestration software market
April 2012	Cancellation Project Olympus	Merging of codebases halted due to CloudStack submission to Apache Foundation	Citrix focus on CloudStack as more mature solution
April 2012	CloudStack submission to Apache Foundation	Facilitate development of open-source cloud orchestration software	Allows variety of stakeholders to compete with VMware’s vSphere

Table 1: Design moves of the OpenStack and CloudStack projects

Appendix: Interview Guide

Our preliminary interviews followed a semi-structured approach. We outlined several general themes (underlined below), each of which consisted of several topics (listed in bullet points below the themes). Depending on the response of the interviewee we asked follow-up questions, either for clarification or if the answer prompted further ideas related to the topic.

General ecosystem issues

- What do you consider to be the key functionalities in the overall cloud architecture (stack)?
- What are the key challenges (both strategic and design related) moving forward?
- What are your expectations of how the cloud computing ecosystem will develop?

Business model and open source development

- Why did you decide to pursue an open source-based development model?
- What are the specific challenges to OSS-based development?
- How does your business model work in terms of value creation and value capture?

Project governance and collaboration

- Can you elaborate on the sequence of decision making in relation to creating or joining OpenStack and/or CloudStack?
- What is your view on Amazon's AWS APIs and how does it affect the development of your project?
- Who do you consider to be your key partners and key competitors?
- Are there any further issues you think that are relevant for our research project?

References

- Adner R. 2006. Match your innovation strategy to your innovation ecosystem. *Harvard Business Review* **84**(4): 98–107.
- Adner, R. and Kapoor, R. 2010. Value creation in innovation ecosystems: how the structure of technological interdependence affects firm performance in new technology generations. *Strategic Management Journal*, 31: 306–333.
- Anderson, P., and Tushman, M. 1990. Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change. *Administrative Science Quarterly*, 35 (1): 604-633.
- Baldwin, C.Y. and Clark, K. 2000. *Design Rules, The Power of Modularity*. MIT Press.
- Baldwin, C.Y. and Henkel, J. 2011. The Impact of Modularity on Intellectual Property and Value Appropriation. Working Paper
- Baldwin CY, Woodard CJ. 2007. Competition in modular clusters. Harvard Business School working paper 08-042. Available at <http://www.hbs.edu/research/pdf/08-042.pdf>.
- Baldwin, C.Y. and Woodard. C.J. 2009. The architecture of platforms: a unified view. In Gawer, A. (ed.) 2009. *Platforms, Markets and Innovation*, Cheltenham, UK and Northampton, MA, US: Edward Elgar, 19–44.
- Besen, S. M., & Farrell, J. 1994. Choosing how to compete: Strategies and tactics in standardization. *Journal of Economic Perspectives*, 8(2), 117–131.
- Borras, M. and Zysman, J. 1998. Globalization with Borders: The Rise of “Wintelism” as the Future of Industrial Competition In Zysman and Schwartz (eds.) *Enlarging Europe: The Industrial Foundations of a New Political Reality*,. University of California Press.
- Boudreau K. 2010. Open platform strategies and innovation: granting access vs. devolving control. *Management Science* **56**(10): 1849–1872.
- Brandenburger, A. and Nalebuff, B.J. 1996. *Co-opetition*, New York: Doubleday.
- Bresnahan T, Greenstein S. 1999. Technological competition and the structure of the computer industry. *Journal of Industrial Economics* **47**(1): 1–40.
- Brusoni, S., Jacobides M.G. and Prencipe, A. 2009. Strategic Dynamics in Industry Architectures and the Challenges of Knowledge Integration. *European Management Review*, 6(4): 209-216.
- Brusoni, S. and Prencipe, A. 2006. Making design rules: a multidomain perspective. *Organ.Sci.*,17(2)179–189.

Cabigiosu, A. and Camuffo, A. 2012. Beyond the mirroring hypothesis - product modularity and inter-organizational relations in the air-conditioning industry, *Organ. Sci.*, 23:672-685.

Casadesus-Masanell R, Yoffie DB. 2007. Wintel: cooperation and conflict. *Management Science* 53(4): 584–598.

Christensen, C. M. 1997. *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business Press.

Christensen CM, Raynor ME. 2003. *The Innovator's Solution: Creating and Sustaining Successful Growth*. Harvard Business School Press: Boston, MA.

Clark, K. B. 1985. The interaction of design hierarchies and market concepts in technological evolution. *Research Policy*, 14(5) 235–251.

Colfer, L.J. and Baldwin, C.Y. 2010. The Mirroring Hypothesis Theory, Evidence and Exceptions, Harvard Business School Working Paper, 10-058.

Eisenhardt, K.M. and Martin, J.A. 2000. Dynamic capabilities What are they? *Strategic Management J.*, 21(10-11) 1105-112

Eisenmann, T.R., Parker, G. and M. Van Alstyne. 2011. Platform Envelopment. *Strategic Management Journal*, 32 (12): 1270–1285.

Farrell, J, and G. Saloner. 1985. Standardization, compatibility and innovation. *Rand Journal of Economics* 16, 70-83.

Garud, R., Jain, S., Kumaraswamy, A., 2002, Institutional Entrepreneurship In The Sponsorship Of Common Technological Standards: The Case Of Sun Microsystems And Java, *Academy of Management Journal* 45 (1), 196-214.

Gawer, A. and Cusumano, M.A. 2002. *Platform Leadership How Intel, Palm, Cisco and Others Drive Industry Innovation*. Harvard Press, Cambridge, MA.

Gawer A, Henderson R. 2007. Platform owner entry and innovation in complementary markets: evidence from Intel. *Journal of Economics and Management Strategy* 16(1): 1–34.

Helfat, C.E. and Peteraf, M.A. 2003. The dynamic resource-based view capability lifecycles, *Strategic Management J.*, 24 997-1010.

Henderson, R. and Clark, K. 1990. Architectural innovation The reconfiguration of existing product technologies and the failure of established firms. *Administrative Sci. Quarterly*, 35 9–30.

Iansiti, M. and Levien R. (2004), *The keystone advantage: What the new dynamics of business ecosystems mean for strategy, innovation, and sustainability*, Harvard Business School Press.

Jacobides, M.G., Knudsen, T. and Augier, M. 2006. Benefiting from Innovation: Value Creation, Value Appropriation and the Role of Industry Architectures. *Research Policy*, 35(8): 1200-1221.

Jacobides MG, Tae CW. 2012. How value migrates within an industry architecture: kingpins, bottlenecks, and evolutionary dynamics. *Academy of Management Annual Meeting Proceedings*. Available at <http://proceedings.aom.org/content/2012/1/1.202>.

Kaplan, S. and Tripsas, M. 2008. Thinking About Technology: Applying a Cognitive Lens to Technical Change. *Research Policy*, 37(5): 790-805.

Katz ML, Shapiro C. 1985. Network externalities, competition, and compatibility. *American Economic Review* 75: 424–440.

Langlois, R.N. and Robertson, R. 1992. Networks and Innovation in a Modular System Lessons from the Microcomputer and Stereo Component Industries. *Research Policy*, 21(4) 297-313.

Liu F, Tong J, Mao J, Bohn R, Messina J, Badger L, Leaf D. 2011. NIST cloud computing reference architecture. *National Institute of Standards and Technology Special Publication 500-292*. Available at http://www.nist.gov/customcf/get_pdf.cfm?pub_id=909505.

MacCormack, A., Rusnak, J. and Baldwin, C.Y. 2012. Exploring the duality between product and Organizational architectures a test of the mirroring hypothesis. *Research Policy* 41 (2012) 1309–1324.

Moore, J. 1996, *The Death of Competition: Leadership and Strategy in the Age of Business Ecosystems*, New York: Harper Collins.

Morris CR, Ferguson CH. 1993. How architecture wins technology wars. *Harvard Business Review* 71(3): 86–96.

Murmann, J. P. and Frenken, K. 2006. Toward a systematic framework for research on dominant designs, technological innovations, and industrial change. *Research Policy*, 35 925–952.

Nelson, R. & Winter, S. 1982. *An evolutionary theory of economic change*. Cambridge, MA: Belknap Press/Harvard University Press.

Parker GG, Van Alstyne MW. 2005. Two-sided network effects: a theory of information product design. *Management Science* 51(10): 1494–1504.

Parker GG, Van Alstyne MW. 2013. Innovation, openness, and platform control. Working paper. Available at <http://ssrn.com/abstract=1079712>.

Pisano, G. P. and Teece, D. J. 2007. How to capture value from innovation: Shaping intellectual property and industry architecture. *California Management Review*, 50: 278-296.

- Richard PJ., Devinney TM. 2005. Modular Strategies: B2B Technology and Architectural Knowledge, *California Management Review* 47(4): 86-113.
- Rochet J-C, Tirole J. 2003. Platform competition in two-sided markets. *Journal of the European Economic Association* 1(4): 990–1029.
- Santos, F. and Eisenhardt, K. 2009. Constructing Markets and Shaping Boundaries: Entrepreneurial Power In Nascent Fields. *Academy Of Management Journal*, 52(4): 643-671.
- Schilling, M.A. 2000. Towards a General Modular Systems Theory and its Application to Inter-firm Product Modularity. *Academy of Management Review*, 25: 312-334.
- Shapiro, C. and Varian, H.R. 1998. *Information Rules: A Strategic Guide to the Network Economy*. Boston, MA: Harvard Business School Press.
- Simon, H. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106 (6) 467-482.
- Slywotzky AJ. 1996. *Value Migration: How to Think Several Moves Ahead of the Competition*. Harvard Business School Press: Boston, MA.
- Suarez, F.F. 2004. Battles for Technological Dominance: An Integrated Framework. *Research Policy*, 33(2): 271-286.
- Teece, D.J. 1986, Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15 (6): 285-305.
- Teece, D.J, Pisano, G. and Shuen, A., 1997. Dynamic capabilities and strategic management. *Strategic Management Journal*, 18 (7): 509-533.
- Tripsas, M. 1997. Unraveling the Process of Creative Destruction: Complementary Assets and Incumbent Survival in the Typesetter Industry. *Strategic Management Journal*, 18: 119-142.
- Ulrich, K. 1995. The Role of Product Architecture in the Manufacturing Firm. *Research Policy*, 24419-44.
- Utterback, J.M. and Abernathy, W.J. 1975. A Dynamic Model of Product and Process Innovation. *Omega*, 3 (6): 639–656.
- West J. 2003. How open is open enough? Melding proprietary and open source platform strategies. *Research Policy* 32(7): 1259–1285.
- Williamson, P.J. and De Meyer, A. 2012. Ecosystem advantage: how to successfully harness the power of partners. *California Management Review*, 55(1): 24-46

Woodard, C.J., Baldwin, C.Y, Tee, R 2012. The Value of Open Standards in Layered Modular Architectures, Working Paper.

Woodard, C. J. and J. West 2009. “Strategic Responses to Standardization: Embrace, Extend or Extinguish?” In G. Cattani et al., eds., *Project-Based Organizing and Strategic Management (Advances in Strategic Management, Vol. 28)*, Emerald, 263–285.

Woodard, C.J., Ramasubbu, N., Tschang, F.T. and V. Sambamurthy 2013. Design Capital and Design Moves: The Logic of Digital Business Strategy. *MISQ Special Issue: Digital Business Strategy: Toward a Next Generation of Insights*, forthcoming.

Yoo Y, Henfridsson O, Lyytinen K. 2010. The new organizing logic of digital innovation: an agenda for information systems research. *Information Systems Research* **21**(4): 724–735.

Zhu F, Iansiti M. 2011. Entry into platform-based markets. *Strategic Management Journal* **33**(1): 88–106.