

Partnering for Value Perfection and Business Sustainability in the Cloud Services Brokerage Market

Richard Shang
University of North Florida
d.shang@unf.edu

Robert J. Kauffman
Copenhagen Business School
rk.digi@cbs.dk

Abstract

The cloud computing services market has advanced in the past ten years. It now includes IT services from fundamental computing to cutting-edge AI capabilities. With the widespread adoption of cloud services, clients are facing the fact that they are utilizing cloud resources at a sub-optimal level. Cloud services brokerage (CSB) emerged in the market to fill the need for cloud resource management and risk mitigation. Based on analysis of the cloud market and the case of CSB and related activities in North America, we offer theoretical analysis for how value creation works, its impacts on the cloud ecosystem, and factors that support the value and business sustainability of the CSB intermediation process.

1. Introduction

In the 2010s till now, cloud computing advanced quickly, and public cloud adoption reached 94%, according to a survey by Flexera [16]. The landscape of cloud services has also expanded from its original focus on IaaS (infrastructure), PaaS (platform), and SaaS (software) to many *X-as-a-service offerings*. They include everything from the Internet of Things (IoT) to machine learning (ML), artificial intelligence (AI), and analytical services on the cloud.

As cloud services users rely more on IaaS and PaaS for data storage to advanced analytics, cloud resource management has become a big challenge. There are variations in how cloud services are provided, charged, and consumed in the market. Even though clients benefit from their affordability, flexibility and scalability, most can use cloud resources better. *Sub-optimal utilization* of cloud resources is caused by there being so many options available: reserved, on-demand, and spot instances of different instance types and sizes; and instances in different geographic zones with varying prices and availability. Moreover, the industry has seen a trend toward *multi-cloud services*, a practice in which enterprises use multiple cloud providers for different application workloads to meet varied technical or work requirements [16].

The optimization of cloud resources from different providers) requires data analytics, machine-based

value maximization, and management, which users are incapable of. McKinsey [22] has referred to the need for cloud management services to free clients from uncertainty about adoption, use, workload management, and demand-and-supply conditions in day-to-day operations as the *progressive cloud*. The new capabilities are based on moving “*from basic storage and networking to innovative offerings like advanced analytics, ML, and virtual-reality development.*”

Cloud services brokers (CSBs, or just *brokers*) are natural intermediaries that create value-added offerings on top of the services of cloud providers. The standard value propositions of digital intermediaries are commonly attributed to reducing transaction costs, aggregating, or matching demand and supply, reducing risks, and addressing trust and moral hazard. As the cloud services ecosystem evolves, we have observed rapid development of CSB intermediaries, increasing the likelihood of further growth in the burgeoning area of cloud business practices.

CSBs have developed platforms to enable clients to manage multiple systems, services and providers. They apply various resource management techniques to provide more streamlined and reliable services to clients, reducing complexity and minimizing costs. The value-creation of CSBs, which is different from that of traditional intermediaries, has evolved quickly. In 2012, for example, Panettieri [39] identified 40 CSB start-ups. By 2016 though, the market no longer recognized the early 2010s’ conceptualization of it as a unique, new business model. The proliferation of CSBs was like the many B2B e-markets that emerged prior to the DotCom crash. Fewer firms have claimed to be cloud aggregators though, and others expanded their services and grew via cross-market platform envelopment. Still others exited the CSB market altogether via acquisition, repositioning their businesses, or failing outright.

More recently, TechRadar [14] listed only the top-five firms, suggesting there was growth in aggregate business revenues, but also increasingly-concentrated competition. There is still an active market though it is different than many expected, and CSB cross-platform envelopment and hybrid cloud services are some of the

next research targets of interest.^{1,2} Orbis Research [38] pointed to the same market leaders, but they are among many other CSB intermediaries now. The CSB model as just a brokerage platform is no longer viable in isolation in the ecosystem. CSBs must create more innovative services with their assets.

We offer a business strategy interpretation of the role of *IT services intermediaries*, which expands the theoretical bases used to understand them. We also explain the economic rationale for why CSBs can develop sustainable intermediary positions in the multi-cloud services value chain. Some key characteristics we point to are: *co-specialization of assets* for value creation; *increasing centrality of information* as an asset; how CSBs' strategic alliances create a basis for *partnering for value perfection* and *market sustainability*; and why transformational *customer and partner informedness* are enabled via data sharing and co-creation of value for clients.

The 35th anniversary of HICSS' SITES Mini-track motivated us to develop this exploratory case study of CSBs to understand the value of services they offer in an increasingly competitive and complex market. This is a hallmark of research on strategy, information, technology, economics, and society during the past four decades, to study how the information-based transformation of strategy and society evolve [8, 19]. Insights from this study will shed light on the value creation and sustainability of intermediaries in other markets [7], as new products and services arise and are affected by digital transformation [5, 20].

2. Theoretical Background

The U.S. National Institute of Standards and Technology (NIST) defines a CSB as “*an entity that manages the use, performance and delivery of cloud services, and negotiates relationships between cloud providers and cloud consumers*” [32]. CSBs combine and integrate cloud services and enhance them by providing specific capabilities and value-added services to cloud consumers. Predictive analytics, ML, and other methods that support value-conscious use of limited server resources have emerged for application in cloud services [12, 13, 33]. Empowered by these resource management techniques, CSBs have grown in the market and provide more reliable services to more clients at a lower cost [45]. This is a strong indication of the greater tradability that CSBs create as digital intermediaries in the market for such services. Without the overhead of managing and maintaining a large pool of infrastructure and software stacks, they

have been able to focus on innovative ways of offering services, and developing a customized portfolio of offerings.

Although CSBs are a new group of intermediaries in cloud markets, they function like intermediaries in other settings with similar attributes. Traditional intermediaries provide services such as price-setting, transaction processing, inventory management, quality guarantees, and service immediacy and monitoring [59]. In e-markets, new forms of IT-enabled intermediation now offer trust provision, aggregation, one-stop shopping, information exchange facilitation, and information filtering [5, 15].

Intermediation is based on three enablers: imitation and the inability to fully appropriate transaction value; ownership of co-specialized assets; and economies of scale [7]. *Co-specialized assets* are the expertise, customer base and relationships of the intermediaries that enable their existence in the value chain. As cloud computing has evolved, selection and management of cloud services have become too complicated for cloud clients to manage alone. CSBs, with their expertise in resource management and optimization, can provide streamlined and reliable services to clients as a result. Also, via their relationships with primary cloud providers, they can function as a single point of entry to manage multiple differentiated cloud services.

CSBs will not succeed if they cannot achieve economies of scale though. Strategic alliances with cloud providers help CSBs to achieve larger customer bases. Moreover, strategic alliances facilitate the development of innovative products and services to address rapid technological changes in the business environment. There are three benefits of strategic alliances in B2B e-markets: their capacity to bring in new resources; identification and endorsement of quality; and flexibility for absorbing uncertainty.

Strategic alliances allow firms access to resources and assets and thus they can build capabilities for developing and offering new products and services or marketing and distributing them to customers. Dai and Kauffman [11] theorized that this process is one of *partnering for perfection*. Through alliances, partnering firms' resource complementarities build collective IT capabilities that support integrated and joint innovation that can enhance their collective profit [44]. Han et al. [24] investigated the *value co-creation of open innovation alliances* (OIAs). They found that the degree of their openness is associated with the level of returns accruing to the partnering firms.

¹ TechRadar's top-5 CSBs in the current market include: AWS, IBM Cloud Brokerage, Cloudmore, Jamcracker CSB, and Dell Boomi.

² It identified more CSB market firms, including four of Tech-

Radar's five, and HPE, Rightscale, Wipro, Arrow, Accenture, Doublehorn, ActivePlatform, Incontinuum, DXC, Cognizant, BitTitan, Nephos, Computenext, Fujitsu, ATOS, and Cloudreach.

Further, by linking with more well-established partners, a firm can enhance perceptions of its product and service quality as well its capabilities in the marketplace. While strategic alliances enable firms to obtain access to assets rapidly, and encourage costs and benefits sharing, they give firms the flexibility in response to changes in demand or other aspects of their business environment. Prior research also found that market uncertainty makes alliances a preferred strategy for product functionality innovations and product promotion for acceptance in early-stage technology development and commercialization [43].

More recently, Gartner [17] indicated that it views a CSB as having a new IT role and business model in the market. Gartner associates this with the provision of PaaS capabilities that add value for clients via resource aggregation, service integration, and service arbitrage. It further distinguishes between *CSBs-as-enablers* that provide technology for cloud services implementation versus *CSBs-as-mixed-services-providers* that sell a blend of technology, people and methods to implement, and manage cloud projects. Katrick [25] defined the essential benefits of CSB intermediaries as: integrated service management for all-year uptime; access to IT resources anytime, anywhere to ensure cloud data access and processing; flexible resource scaling for optimizing spend as a firm's business changes; lower total cost of ownership to mitigate new capital cost requirements; and automated self-service delivery to increase agility in cloud app deployment.

With input from participants in MIT's 2018 Platform Strategy Summit, Parker et al. [41] pointed to the more general aspects of how platform strategies in markets for different kinds of services create value. One is that platforms provide the technology, scale and risk management capabilities to create the basis for intelligent risk absorption. The growth of platforms toward specialization and the use of advanced decision engines make it possible to address value creation in highly sophisticated ways. The authors noted the Uber platform's use of five million forecasts every minute to balance ridesharing supply and demand, for example. Platforms are also a basis for the discovery of new revenue streams that are based on data-driven innovation and tech-savvy data analytics talent. In addition, platform strategies enable firms to stress how to build value through the diversity of racial, ethnic, educational of the teams that handle them. Finally, fragmentation of markets – cloud services among them – provides a fertile context for synergy and

sinnovation.

Thus, it is important for CSBs to connect with the high-performance possibilities of more well-developed platform markets.³ For matching vendors and clients that transact with each other using system resources, platforms focus on optimizing the value created by the ecosystem to maximize growth [40]. In the cloud market, platforms like AWS have promoted value creation by implementing flexible provision and pricing structures. For cloud platforms, development of mechanisms that enable flexible, rapid, and seamless integration of services by CSBs is important also.

These capabilities are important to the success of CSBs in the market. We adopted the general theoretical background to guide analysis of the role of CSBs with these questions in mind:

- How do they achieve sustainable value chain positions with their co-specialized assets?
- How do strategic alliances help them to create the basis for appropriating the value of cloud services? Can firms partner for value perfection?
- Are we seeing the emergence of cloud service analytics and management platforms that will transform the sophistication of the cloud market for clients that use multi-cloud services?

3. Research Context

Enterprises are moving their data processing workloads to the public cloud to realize the benefits of such services for improved agility and performance. Meanwhile, they are faced with increasing complexity and higher than expected costs. Navigating this complexity and achieving better ROI are key concerns for enterprise decision-makers as a result.

The complexity of cloud resources lies in the various options available in the cloud services market. For example, Amazon Elastic Compute Cloud (EC2) offers a wide selection of instance types that are appropriate to fit different use cases. Amazon provides five families of instance types classified according to their use cases. They comprise varying combinations of CPU, memory, storage and networking capacity to give the client flexible choice of the appropriate mix of resources for their applications. (See Table 1.) Each instance type includes one or more instance sizes, allowing clients to scale resources to their target workload requirements. The instances are offered with different purchase options, contract lengths, and payment plans in different zones. With flexibility to

³ Our case study partner, Spot by NetApp (Spot.io), was acquired by cloud services market data services leader, NetApp, in 2020 [36]. Together, they are creating “an application-driven infrastructure to enable customers to deploy more applications to public clouds

faster with Spot's ‘as-a-service platform’ for the continuous optimization of compute and storage for traditional IT buyers with enterprise applications, cloud-native workloads, and data lakes.” This lends credence to the argument we made on platform value.

choose the right mix of resources comes the complexity to optimize use of resources available in the market though.

Cloud resources tend to be distributed. For example, a computational task can be completed with one powerful extra-large instance, or with 100 small instances via parallel and distributed computing. An x-large instance with 128 CPUs is much more costly though than the total of 128 nano instances each with 1 CPU. Utilizing different types of instances to minimize cloud spending is difficult for most clients.

Table 1. Options for AWS compute instances

Instance family ^(a)	General-purpose (9 instance types) Compute-optimized (3 types) Memory-optimized (8 types) Accelerated-computing (5 types) Storage-optimized (4 types)
Instance size	Nano, micro, S / M / L / XL, metal
Purchase option	On-demand; reserved; spot
Contract length	None; 1 year; 3 years
Offering class	Standard; convertible
Payment plan	Upfront (no / partial / full)
Locations	HK; Mumbai; Osaka, Tokyo; Seoul; Singapore; Sydney; Toronto; Frankfurt; Dublin, London; Paris; Stockholm; Sao Paulo; and VA, OH, CA, OR (USA)
Oper. sys. (OS).	Linux; RHEL; SUSE; Windows

Note: ^(a) Each type in a family offers different features suited to jobs that require different processing (e.g., burst vs. continuous), CPU cycles (e.g., high performance vs. standard processors), and memory (stochastic load vs. all data-in-memory processing [2]).

Things get more complicated when users consider different purchase, contract and offering options, as well as when they can choose instances from different locations. Reserved instances are less costly than on-demand instances when users are willing to lock in contracts for a period of one or three years. Users can get extra discounts when they make payments up-front instead of on a *pay-as-you-go* basis. Further, prices of instances vary at different locations. U.S. clients can use compute instances at any of five locations in North America, for example, without any difference in the quality of services, but at different prices.

Often, a client will utilize different instance types for different use cases. So, there is a need for the flexibility to change instance types when a client is locked into a contract of reserved instances. AWS *convertible reserved instances* offer clients the option to exchange one instance for another convertible reserved instance with a configuration – which makes this into something like a financial derivative. This may include the instance family, OS, and tenancy, as long as the target instance is of equal or higher value. Thus, even for clients who use only reserved instances (with predictable costs), to optimize cloud resource costs, vendors must have a portfolio of different types.

To ensure the scalability of services, cloud services providers must always have unused capacity ready for

immediate growth. They do not want idle resources that add to overhead costs. To utilize their spare capacity, in 2009, Amazon announced a spot instance offering [1]. In terms of hardware and technical parameters, spot instances are exactly the same as regular instances: only the pricing and availability are different. Spot instances, which can be acquired for large discounts when compared to regular on-demand prices come with a major caveat. Whenever spare capacity is needed for standard compute processing, the spot instances can be terminated, with best-effort, two-minute notice. AWS created a market based on supply and demand causing the price for compute instances to change as use and spare capacity fluctuate.

Savings in the marketplace can reach up to 90% compared to on-demand services, but with risk of termination [3]. Following AWS, Google launched its own spot services under the name “Preemptible Virtual Machines” in 2016 [29]. Google makes its *virtual machine* (VM) services available at fixed prices, but they can only run for up to 24 hours. Preemptible VM's can cost up to 80% less than regular instances. Microsoft Azure launched its own spot services under the name “Low-Priority VMs” in 2017, also at fixed prices [34]. There is no time limit with Azure Low-Priority VMs, BUT when capacity is needed, Microsoft can revoke these instances at any time, with no warning.

Cloud services have opened up new opportunities for integrating business processes across firm boundaries. There is an emerging trend with container-based application development alongside multi-cloud strategy adoption. According to Gartner, for example, “by 2022 more than 75% of global organizations will be running containerized applications in production, which is a significant increase from fewer than 30% today” [42]. In container-based development, abstraction from the host operating system makes containerized apps portable and able to run uniformly and consistently across any cloud. New cloud applications can be built from the ground up as containerized microservices, breaking a complex application into a series of smaller specialized services. Tools like Docker are built around container engines where containers act as portable means to package applications [37]. CSBs are in a position in the cloud value chain with their capabilities to create innovative value-added services. Multi-cloud development presents an opportunity to capture market shares and bring in disruptive ITs.

4. Case Analyses on the CSB Strategies

Our analyses represent a case study of CSB intermediation. This approach allowed us to gain an

in-depth understanding of cloud services intermediation strategy in the cloud services ecosystem and the related management issues [28]. Our research involved the use of both primary sources (product demonstrations presented by tech professionals at a focal firm) and extensive secondary sources (e.g., use cases, technical reports, whitepapers, and press releases). Through the demos and technical reports, we obtained an understanding of the technical underpinnings of CSB intermediation. Then, we analyzed the use cases to examine how a CSB can achieve a sustainable position in the value chain of cloud services.

4.1. Value Chain with Co-specialized Assets

One of the focal firms, Spot, is a start-up funded in 2015 in Tel Aviv, with offices after its acquisition by NetApp in San Francisco, New York, London, and elsewhere. It provides cloud services aggregation, arbitrage, and technical intermediation, with a key focus on cloud services automation and optimization.

Firms reported in Spot's use cases heavily rely on IT infrastructure to support their business operations. They offer IT services consulting or SaaS, host platforms or websites for customers, or conduct computation-intensive analytical tasks. The firms understand the benefits of using cloud services and have migrated part or all of their operations to the cloud. Several are "cloud native" companies that have built services and establish platforms specifically on top of cloud-native architectures. This agile approach has allowed them to improve their ability to quickly respond to change and reduce operation costs. Despite the various instance types and flexible pricing, they realized that running in the cloud was not enough for product agility and reduced costs. As Chegg.com, a U.S. textbook rental firm, suggested: *"as soon as we unlocked the gates, the number of services ... launched in ECS [Amazon Elastic Container Service] skyrocketed, which was great for [e]ngineer[ing] productivity, but we needed to ... manage our growing [operating expenses]"* [50].

Changes in business requirements and associated requests for cloud resource can be very dynamic. An example is one of Spot's clients, GumGum, a U.S.-based applied computer vision firm. It *"received around 4 million requests per minute at low traffic points. For peaks, this number grew to over 18 million"* [51]. To handle the dynamic changes, firms would have been burdened with the overhead of managing the launch or termination of instances constantly while trying to maintain the same quality of service.

Different workloads require different computing resources. Some are optimized for memory, CPU, or

networking – or a combination of all three. One way to reduce costs is by combining different instance types and sizes to make better use of cloud resources. Ensuring cost optimization is a difficult task when workloads cut across different environments and cloud providers though. Each environment can have unique usage, compute demand, configuration, and integration needs, so optimization strategies need to be versatile and easy to scale. An example of this is Duolingo, a U.S.-based language education platform. They *"saw their costs rise and made an attempt to reduce [them] by combining the use of reserved instances and EC2 spot instances. As time went on, Duolingo realized they [were] not using reserved instances effectively, because they were not able to prioritize using them before ... spot instances. [T]hey were left with unutilized reserved instances which is money wasted"* [52].

Spot instances offer savings, but they come at the risk of downtime if not managed properly. With up to 90% savings across providers, being able to take advantage of the savings that spot instances and spare capacity have to offer can be the key to cutting cloud costs. However, much like financial market trading, it requires advanced algorithms to predict changing prices and how much processing capacity is available at any given time. While termination is unavoidable with spot instances, sophisticated workload management is required to enable load balancing and backup. A client that uses spot instances faces the risk of breaching the uptime *service level agreement* (SLA), impacting its business if any spot terminations leave applications unavailable. An example is Zalando, a German online fashion retailer. In deciding to leverage spot instances, they found that *"the 2-minute notifications on spot termination [was] not enough for the vast majority of their applications"*, and *"their critical services [became] unavailable or [had] degradation of performance"* [53]. Another is Wiser, a U.S. analytics firm. It learned that *"spot instances were too risky for us, even though they were super cheap. The main issue was their unreliable availability. There were days we lost 90% of our Spot capacity in a split second"* [54].

Spot's position in the cloud ecosystem lies in its capability with co-specialized assets to address the above challenges. The firm's application scaling service, Elastigroup, uses predictive algorithms to leverage the spot instances of major cloud vendors to reduce the risk and complexity. Elastigroup *"predicts the interruption of spare compute instances and seamlessly replaces them, ensuring that the application is always available. [It] monitors spare compute prices to locate your application in the most cost-efficient market and launches on-demand*

instances only when no spare capacity is available. [And] you can take advantage of up to 80% in cost savings while maintaining 100% availability, backed by an SLA” [55]. For clients that have purchased reserved instances, whenever Spot identifies unused reserved instances, it moves relevant workloads from spot to reserved instances. This helps clients use what they paid for instead of incurring costs for unneeded instances.

We attribute Spot’s ability to automate and optimize cloud resources to its analytics capability. Further, the *financification of cloud computing services* [26], emergence of data analytics, and innovative managerial processes from cloud intermediaries have caused the information-driven transformation of cloud services strategy. It now operates like a supply and demand-based financial market. This requires clients to make fast, computer-aided decisions to optimize cloud expenditures in a complex and competitive market.⁴

Past research has examined the volatility of spot prices and developed ways to predict them to reduce client risk and improve spot market services utilization [6, 27]. Building on past work, Spot’s Elastigroup analyzes the spot instances market using machine learning and predicts interruptions ahead of time. When its algorithms identify any spot instances at risk for termination, prior to that it launches new instances to replace them. Consider the evidence from Freshworks, a U.S.-based SaaS vendor: *“In most cases, Spot was able to predict an interruption 15 minutes prior to AWS’ notification and immediately scheduled the EC2 Spot Instance for replacement, and in edge cases, we were even notified 20 minutes in advance”* [56]. This frees clients from monitoring and bidding in the spot market, determining the best availability, and manually replacing interrupted instances.⁵

Many firms use more than one cloud service vendor to cut costs, prevent vendor lock-in, improve disaster recovery, and make migration easier for some data and applications. A 2019 Gartner survey of public cloud users noted that 81% are working with two or more vendors [18]. Multi-cloud is a differentiator for CSBs. Spot’s Elastigroup supports applications run on spare capacity from multiple cloud providers: AWS, Azure, Google Cloud, and Alibaba Cloud. An example of this is the case of Samsung SDS: *“[The] more applications, environments and providers Elastigroup can support, the higher the level of optimization” which Samsung’s customers can achieve* [58].

In the market for cloud-native analytic database services, another focal firm that we studied, Snowflake, is up against the cloud providers and driving customers to its platform. Snowflake, founded in 2012, is a data warehouse distributed across availability zones of the platforms (AWS, Microsoft Azure, or Google Cloud) on which it runs. It is designed to operate continuously and tolerate component and network failures with minimal impact to customers. Due to security and costs, customers want to keep data in the cloud. Customers whose data are stored in multiple clouds demand low-latency access to that data though.

Snowflake creates value by automating its portion of the data flow by abstracting complexity related to data location and optimizing its portion of the stack to get insights irrespective of data location. For example, Happy Money, a consumer financial products provider based in California, found that *“With Snowflake, Happy Money, doesn’t have to worry about infrastructure, maintenance, or capacity upgrades. Instead, the team can focus on providing business value”* [47]. Similarly, Big Fish Games, a developer and publisher of game franchises, concluded that *“Snowflake allows us to not be locked within [Google’s Cloud Platform] GCP, and allows us to collaborate with diverse partners who may work in different public clouds”* [46].

4.2. Partnering for Value Perfection: Strategic Alliances with Cloud Providers

As cloud usage grows and environments become more complex, keeping pace with changes in the cloud market can seem impossible. For example, AWS provides a large selection of options to fit different use cases, and clients can mass-customize combinations of specifications. However, every client’s use case is unique, and personalizing services to suit them according to their profiles is beyond Amazon’s reach given the size of its market. While Amazon focuses on broadening and improving tools, processes and infrastructure, AWS solution providers ease this burden by playing a critical role in helping to properly manage and scale cloud use cases.

The AWS Solution Providers Program creates opportunities for cloud service partners that go beyond reselling, to providing value-added services that help clients to be successful in the public cloud. Amazon incentivizes partners with tiered discounts that reward them for their technical investments. In addition, the program enables authorized solution providers to

⁴ The developments in cloud services are what Lewis [30] described as real-time data and optimization of high-frequency market trading.

⁵ *Edge cases* occur when computing resources are closer to the

client, not in a large cloud data center distant to the user. With close distance to the computing, the less the notification time is.

create, manage and bill customer accounts end-to-end. This is important for Spot to reduce transaction costs and enable seamless integration as customers switch to its platform. With a simple download from Spot's console, clients are able to import existing AWS workloads and have them run on spot, reserved and on-demand instances. FindHotel, a Dutch travel firm, for example, *"migrated their largest, high availability system from Fargate to Elastigroup in just a few short days and very quickly achieved massive cost reduction on their cloud-compute spend"* [57]. Alliances between cloud providers and CSBs enable quick responses to changes and create new services.

Snowflake runs on AWS, Azure and GCP, and is a major partner of the cloud providers. The cloud providers all have products (AWS Redshift, Google Big Query and Microsoft Azure) that compete with Snowflake and want customers to use their database services. Also, they need Snowflake to run in their clouds because it sells infrastructure services, expands services optionality for customers, and supports the ongoing evolution of the ecosystem.

Cloud services clients are looking for proven expertise as they evaluate requirements for taking steps with adoption. Amazon provides benefits and support to partners who are most adept at providing the services and capabilities that clients are demanding. This helps the partners attract new customers while expanding existing environments via new add-on services. For example, AWS endorses CSBs with placements in its AWS Partner Solutions Finder, listings with the Designated AWS Solution Provider or AWS Competency Logos, and involvement in press releases and marketing support. The reinforcements signal solution providers' expertise as being officially certified by AWS and having a proven track record in managing cloud services across scenarios. For example, the AWS Competency Program *"is designed to identify, validate, and promote AWS Partners with demonstrated AWS technical expertise and proven customer success"* [4]. The AWS Competency status differentiates Snowflake as an Amazon Partner Network (APN) member that delivers highly specialized technical proficiency. It enables providers to build opportunities by offering differentiated services like cloud migration.

5. Sustainability of CSB Intermediation

The value created by CSBs arises in a way that is similar to how wealth management services brokers support the investment needs of high net-worth clients on behalf of multiple financial institutions. Each may have limited services or delivery capabilities to

address the full spectrum of their clients' needs. But what makes intermediaries such as Spot and Snowflake sustainably beneficial for clients over time? We next discuss insights drawn from our field work and theory.

5.1. Co-specialized Assets

An intermediary will be stable and sustainable when the value it offers clients persists for product and service transactions that result in profits that clients cannot achieve alone [11, 59]. As we have pointed out, firms wishing to utilize cloud resources found reserved instances too inflexible and spot instances too difficult to manage. Without the capability to effectively and efficiently reduce the risk of using spot instances, most clients' cloud market spending would be sub-optimal.

As CSBs provide services on top of cloud providers' resources, they have to embrace specializations that create added value for customers. Brokers like Spot specialize in services intermediation and leverage predictive algorithms to assist in automating cloud workload automation and cost optimization. Others may provide services aggregation. For example, ComputeNext, another U.S. cloud brokerage platform provider, helps cloud services clients by enhancing their ability to effectively manage their accounts and services through its "Cloud Kiosk." Clients can search, discover, and procure cloud infrastructure in a multi-cloud market, which *"enables end-to-end transactions across [its] platform"* and promotes *"choice and efficient computing through federation"* [9].

Federation in cloud computing results in the possibility of generating benefits through a partnered process of value co-creation [11]. This can happen in two ways. One is between the CSB intermediary and its clients, where joint data processing and problem-solving activities can be pursued – beyond the basic value proposition of connected use of computing processes to reduce costs in readily identifiable ways. End-to-end, platform-wide transaction-making offers very high opportunities for operational ROI via the cloud, and a CSB's key role as a value-creating intermediary will work to strengthen its position as a central player in the cloud services value chain.

5.2. Data as an Economic Asset

Cloud services rely on real-time data access and high-power transaction processing. Such data are an economic asset for firms involved in sophisticated intermediation services. For example, Refinitiv, a financial market data and infrastructure provider, has described a *"real-time data offering in AWS [that] provides access to information on more than 80 million financial instruments, including exclusive data*

sets, to a range of businesses that include banks, brokers, asset managers, hedge funds, payments and remittance companies or ... any firm looking to track financial markets intraday” [10].

Spot’s value-creation is achieved through data analytics and predictive technology, as noted earlier. The usefulness of prediction is enhanced when large-scale historical datasets are available. This is where Spot has the advantage with observations on spot instance availability and performance datasets over a period of years. Based on historical and statistical data, its Elastigroup predicts interruptions ahead of time and automatically migrates instances into different instance types, zones, or falls back to on-demand instances to ensure high availability and consistency.

As firms use analytics to create value, it needs to be recognized that data are perhaps the *most valuable asset in modern business*, along with the informedness that they create for producers and consumers of services. This requires practices that consistently collect, manage and monitor data efficiently throughout the enterprise, and make it available for effective sharing. For example, Petco, *“used data sets from Snowflake Data Marketplace to understand and better predict COVID-19’s impact on stores. Data was behind their decision to launch curbside delivery, same-day delivery, and other capabilities” [48].*

5.3. Strategic Alliances

Strategic alliances with cloud providers are essential to CSB value creation. They enable CSBs to act as a point of contact for multiple cloud providers – managing relationships with different providers for customers and providing a single source for billing, reporting and support. Through alliances, CSBs can resell and integrate services at competitive (less than retail) prices due to the volume they can offer to providers. As enterprises adopt multi-cloud strategies, CSBs can offer services online with solutions that provide flexibility to choose various providers.

Strategic alliances allow CSBs to complement cloud services. Clients want to monetize their cloud investments for personalized enterprise-class services while keeping costs down, so CSBs customize cloud services and identify the best-in-class solution by project; integrate cloud resources to provide added-value services; enforce or enhance SLAs; and customize user interfaces and management consoles.

Alliances between CSBs also enable co-creation of value-added services with customers. Among Spot’s use cases, five firms operate cloud consultancies or managed service providers. Through partnerships, they create benefits by leveraging Spot’s capability with cloud cost optimization. The alliances enable them to implement new solutions and features to

satisfy their customers’ needs and foster their loyalty.

5.4. Enhanced Informedness

A sustainable CSB intermediary will implement strategies for providing services to allow clients, suppliers and vendors to be effective in delivering the right products and services – based on the buyer, supplier and intermediary informedness they co-create [31]. Enterprises need visibility into cloud resources so their cloud workloads can be monitored in a transparent way and configured to improve performance or reduce costs. Better understanding of cloud computing services consumption helps to identify optimization opportunities and allocate resources.

Spot’s platform provides full visibility into cloud costs and usage broken down by services, applications and tasks, enabling them to create budgets based on accurate usage and projections. This makes it a *data-as-a-service* (DaaS) provider, in keeping with our observations on *data as an economic asset*. Returning to our prior example for support for this, FindHotel *“gained clarity into cost allocation using Elastigroup’s visualization of their EC2 spend broken down to the container level. This provided better governance and allowed them to ascertain ROI for business operations” [57].* Also, Snowflake’s solutions have helped Vix Technology, a UK provider of public transport ticketing systems, *“to understand what is going on in real time within their transport networks. They have access to a series of customized dashboards that allow them to slice and dice their data in any way they wish using whatever tools they choose” [49].* Overall, the evidence supports our theory.

This demonstrates the *economic value of informedness* in cloud computing. In hospitality, the sophistication of revenue management, room pricing, price-bucket adjustments, and revenue and demand consolidation requires cloud data processing based on high-end computing. Dynamic room pricing and supply-and-demand management in hospitality – with increasingly sophisticated revenue yield algorithms – are closely aligned with what we observe in cloud services too. So, it is natural to recognize that IS specialists for such large firms are looking for CSB support that creates value and mitigates risk via informedness.

6. Conclusion

This industry case study on CSB intermediation business sustainability contributes new knowledge, especially with a theory of cloud intermediation and the perfection of value co-creating services. This is a

unique aspect of our research, in line with the SITES 35th anniversary mini-track at HICSS. Our work offers new thinking for cloud firms related to strategy and partnering approaches, the services portfolio, and capacity to capture sustainable value in the cloud ecosystem as 100% digital intermediaries. We leveraged cooperation with an innovative CSB intermediary to understand the CSB market and firms within it, as well as the nature of the client relationships in which sustainable value co-creation seems to have been occurring related to the supply of multi-cloud services. This is what makes CSB intermediation uniquely valuable for this exploratory research: multi-cloud services is a game-changer for the CSB ecosystem. More firms will become *mixed-services platform CSBs*.

On one side, the multi-cloud environment increases user demand for CSBs' cloud management and automation, which results in cross-platform effects of attracting cloud providers to partner with CSBs. On the other, the multi-cloud nature of their business enables CSBs to flexibly choose and switch between cloud providers and develop cost-efficient and value-added services on top of the cloud resources that attract more clients. All these put CSBs in a more advantageous position in the cloud services ecosystem. This is an untouched area for business school researchers, although technical researchers have studied some CSB issues [e.g., 21, 24], but still there are many strategic, managerial, organizational, decision-making and economic issues to be examined. Our study builds on and extend previous work on tech intermediation, too. As digital transformation across industries challenges the viability of intermediaries, our study emphasizes the importance of co-specialized assets and alliances for achieving synergies and innovation.

CSB intermediation has rapidly changed from very little activity in 2009 when Gartner defined its three core service areas. Cloud computing seemed like a novelty not a market commodity, as it is today. The global market in 2019 was estimated at US\$5.71 billion, with a compound annual growth rate (CAGR) of 16.6% expected for the period 2020-2025 [35]. A barrier to rapid expansion is how client data and information problems are solved – a growth vector for CSB value, and likely to be folding into their portfolio of services – another value-perfection play.

There are several possible. We studied the quickly evolving cloud computing sector. Instead of a case study approach at the firm level, it may have been

useful to quantitatively analyze multiple firm-to-firm alliances in CSB intermediation – thus, shifting to a different unit of analysis – and to focus on CSB sustainability and platform envelopment from that perspective. Second, it has been fortunate to be studying CSB intermediation when COVID-19 created many barriers for the discovery of new knowledge in our research.⁶ Yet, it may have been more revealing to focus on a data analytics study, by client for the focal firm, with somewhat different research questions on the business value of specific and similar services offered to different clients. The limitations need to be understood for what we have done – not what we could have done though. We hope to take this forward as theory-building work and as research-in-translation for practice.

Acknowledgments. We benefited from access to Spot by Netapp (Spot.io). We also wish to thank Eric Clemons, Qizhi Dai, Thomas Jensen, Helmut Krcmar, Dorothy Leidner, Dan Ma, Somnath Majumdar, Raghava Mukkamala, Jose Parra-Moreno, Hugh Watson, Thomas Weber, and Yiping Yang for helpful discussions, and suggestions on this and related research. The shortcomings are ours.

7. References

- [1] AWS. Announcing EC2 spot instances. Seattle, WA, December 14, 2009.
- [2] —. Amazon EC2 instance types. 2020a.
- [3] —. Amazon EC2 spot instances: Run fault-tolerant workloads for up to 90% off. 2020b.
- [4] —. AWS Competency Programs: Highlight your AWS technical expertise and specialization to differentiate your business. 2020c.
- [5] J.P. Bailey, J.Y. Bakos. An exploratory study of the emerging role of electronic intermediaries. *Intl. J. Elec. Comm.*, 1(3), 1997, 7-20.
- [6] H.K. Cheng, Z. Li, A. Naranjo. Cloud computing spot pricing dynamics: Latency and limits to arbitrage. *Info. Sys. Res.*, 27(1), 2016, 145–165.
- [7] A. Chircu, R.J. Kaufman. Reintermediation strategies in business-to-business electronic commerce. *Intl. J. Elec. Comm.*, 4(4), 2000, 7-42.
- [8] E. Clemons, R. Dewan, R.J. Kauffman, T.A. Weber. Understanding the information-based transformation of strategy and society. *J. MIS.*, 34(2), 2017, 425-456.
- [9] ComputeNext. About us. Redmond, WA, 2020.
- [10] B. Cracknell. Real-time market data cloud technology. Perspectives, Refinitiv, Beijing, August 23, 2019.
- [11] Q. Dai, R.J. Kauffman. Partnering for perfection: An economics perspective on B2B electronic market strategic alliances. In K. Tomak (ed.), *Advances in the Economics of IS*, Hershey, PA: IGI, 2005, pp. 43-79.
- [12] S. Das, A. Du, R. Gopal, R. Ramesh. Risk management and optimal pricing in online storage grids. *Info. Sys. Res.* 22(4), 2011, 756-773.

⁶ This hampered our efforts with conducting additional interviews but was not a fatal flaw for access to key informants. The alternative involved third-party news and information sources, which enabled

us to do more triangulation to get a clearer understanding of what was going on in the market with CSB firms and their clients, and the CSBs' efforts to co-create value through intermediation.

- [13] S. Ding, Y. Li, D. Wu, Y. Zhang, S. Yang. Time-aware cloud service recommendation using similarity-enhanced collaborative filtering and ARIMA model. *Dec. Supp. Sys.* 107, 2018, 103–115.
- [14] N. Drake, B. Turner. Best cloud brokers of 2020: Integrated services for hybrid clouds. *TechRadar*, December 14, 2019.
- [15] O. El Sawy, A. Malhotra, S. Gosain, K. Young. T-intensive value innovation in the electronic economy. *MIS Qtrly.* 23(3), 1999, 303-335.
- [16] Flexera. 2019 state of the cloud report: See the latest cloud trends. Report, Itasca, IL, 2019.
- [17] Gartner. Cloud service brokerage. Stamford, 2020.
- [18] L. Goasduff. Most organizations choose to work with multiple cloud providers, for a host of different reasons. *Smarter with Gartner*, Stamford, May 7, 2019.
- [19] P. Gomber, R.J. Kauffman, C. Parker, B.W. Weber. On the Fintech Revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. *J. MIS.* 35(1), 2018, 220-265.
- [20] N.F. Granados, R.J. Kauffman, H.C. Lai, H.C. Lin. Decommoditization, resonance marketing, and IT: An empirical study of air travel services amid channel conflict. *J. MIS.* 28(2), 2011, 39-74.
- [21] N. Grosov, R. Buyya. Inter-cloud architectures and application brokering: Taxonomy and survey. *Softw. Prac. Exp.* 44(3), 2014, 369-390.
- [22] M. Gu, K. Krishnakanthan, A. Mohanrangan, B. Smolinski. The progressive cloud: A new approach to migration. McKinsey, New York, December 2018.
- [23] M. Guzek, A. Griewek, P. Bouvry, J. Musial, J. Blaze-wicz. Cloud brokering: Current practices and challenges. *IEEE Cloud Comp.* 2, 2015, 40-47.
- [24] K. Han, W. Oh, K. Im, R. Chang, H. Oh, A. Pinsonneault. Value co-creation and wealth spillover in OIAs. *MIS Qtrly.* 36(1), 2012, 291-325.
- [25] J. Katrick. Cloud broker: Essentials for a cloud services brokerage platform. *Medium*, May 16, 2018.
- [26] R.J. Kauffman, D. Ma, R. Shang, J. Huang, Y. Yang. On the financification of cloud computing: Pricing and service delivery mechanism design. *Intl. J. Cloud Comp.*, 2(1), 2014, 114.
- [27] V. Khndelwal, A. Chaturvedi, C. Gupta. Amazon EC2 spot price prediction using regression random forests. *IEEE Trans. Cloud Comp.* 8(1), 2017, 59-72.
- [28] H.K. Klein, M.D. Myers. Set of principles for conducting and evaluating interpretative field studies in IS. *MIS Qtrly.* 23(1), 1999, 67-94.
- [29] F. Lardinois. Google platform's preemptible VMs are now up to 33% cheaper. *Techcrunch*, August 9, 2016.
- [30] Lewis, M. *The Flash Boys: A Wall Street Revolt*, Norton, New York, 2014.
- [31] T. Li, R.J. Kauffman, E. van Heck, P. Vervest, B. Dellaert. Consumer informedness and firm information strategy. *Info. Sys. Res.* 25(2), 2014, 345-363.
- [32] F. Liu, J. Tong, J. Mao, R. Bohn, J. Messina, L. Badger, D. Leaf. NIST cloud computing reference architecture. U.S. Dept. of Comm., Wash., DC, 2011.
- [33] M. Mazzucco, M. Dumas. Achieving performance and availability guarantees with spot instances. In *Proc. 13th Intl. Conf. High Perf. Comp. Comm.*, IEEE Comp. Soc. Press, Wash., DC, 2011, pp. 296-303.
- [34] M. McCrory. Announcing low-priority VMs on scale sets now in the public view. *Microsoft Azure*, Redmond, May 3, 2018.
- [35] Mordor Intelligence. Cloud services brokerage market: Growth, trends, and forecast (2020-2025). Hyderabad, January 2020.
- [36] NetApp. NetApp to acquire Spot. Press release, Sunnyvale, CA, June 3, 2020.
- [37] Newman, S. *Building Microservices: Designing Fine-grained Systems*, O'Reilly Media, San Fran., 2015.
- [38] Orbis Research. Cloud services market and CSB market research. Dallas, February 29, 2020.
- [39] J. Panettieri. Cloud services brokerage company list. *Channel e2e*, Boston, July 1, 2016.
- [40] G. Parker, M. van Alstyne, X. Jiang. Platform ecosystems: How developers invert the firm. *MIS Qtrly.* 41(1), 2017, 255-266.
- [41] G. Parker, M. van Alstyne, P. Evans. How platform strategies continue to create value. *Sloan Mgmt. Rev.* September 24, 2018.
- [42] C. Pettey. Six best practices for creating a container platform strategy. *Gartner*, Stamford, April 23, 2019.
- [43] E. Roberts, W. Liu. Ally or acquire? How technology leaders decide. *Sloan Mgmt. Rev.* 43(1), 2001, 26-34.
- [44] S. Sarker, S. Sarker, A. Sahaym, N. Björn-Andersen, "Exploring value co-creation in relationships between an ERP vendor and its partners: A revelatory case study. *MIS Qtrly.* 36(1), 2012, 317-338.
- [45] R. Shang, R.J., Kauffman, J. Huang, Y. Yang. Client risk informedness in brokered cloud services: An experimental pricing study. *Elec. Comm. Res. Appl.* 39, 2020, 100893.
- [46] Snowflake. Big Fish Games launches and develops games faster with Snowflake. Bozeman, MT, 2021a.
- [47] __. Happy Money transforms consumer lending with performance gains on Snowflake's data cloud, 2021b.
- [48] __. Petco modernizes its data environment for scalable retail analytics with Snowflake. 2021c.
- [49] __. Vix Technology offers real-time access to customer data with Snowflake. 2021d.
- [50] Spot. Chegg cuts EC2 costs by 70% while simplifying ECS infrastructure management. San Fran., 2020a.
- [51] __. GumGum utilizes [Spot] to help drastically reduce costs without compromising availability. 2020b.
- [52] __. How Duolingo reduced ECS costs by 65%. 2020c.
- [53] __. Zalando runs mission-critical workloads on spot instances with [Spot]. 2020d.
- [54] __. [Spot] case study – Wise. 2020e.
- [55] __. Introducing Elastigroup. 2020f.
- [56] __. Freshworks integrates [Spot] Elastigroup to automate AWS-Opsworks on spot instances. 2020g.
- [57] __. FindHotel streamlines ECS workload automation while reducing costs with [Spot] Elastigroup. 2020h.
- [58] __. Samsung SDS optimizes AWS and Azure costs with automated compute provisioning. 2020i.
- [59] D.F. Spulber. Market microstructure and intermediation. *J. Econ. Persp.*, 10(3), 1996, 135-152.