



Measuring the Business Value of Cloud Computing: Emerging Paradigms and Future Directions for Research

Theo Lynn, Pierangelo Rosati, and Grace Fox

Abstract Much of the research on measuring the business value of cloud computing examines cloud computing from the perspective of a centralised commodity-based aggregated conceptualisation of cloud computing, largely based on the NIST reference architecture. Advances in new processor architectures and virtualisation combined with the rise of the Internet of Things are not only changing cloud computing but introducing new computing paradigms from the cloud to the edge. These new paradigms present both opportunities and challenges, not least managing complexity several orders of magnitude greater than today. Yet, academic research on measuring the business value of cloud computing is lagging practice and remains far removed from these innovations. New research is required that explores the relationship between investments in new cloud computing

T. Lynn (✉) • P. Rosati • G. Fox
Irish Institute of Digital Business, DCU Business School, Dublin, Ireland
e-mail: theo.lynn@dcu.ie

paradigms and business value, and the measurement thereof. This chapter explores a selection of these new paradigms, which may provide fruitful research pathways in the future.

Keywords Function-as-a-Service • Serverless computing • Containerisation • High performance computing

7.1 INTRODUCTION

Our interest in measuring the business value of cloud computing stemmed from a 2017 multi-disciplinary survey of the literature we conducted (Rosati et al. 2017). Our findings at that time highlighted a number of worrying issues in the 53 papers published from 2009 to 2016 that we examined. Firstly, the overwhelming majority of studies, in both information systems (IS) and computer science (CS), focussed primarily on one service model, Infrastructure-as-a-Service (IaaS). This is unsurprising as it allows an easier comparison with traditional on-premise computing. In IS, there was also a tendency to conflate all service models as “the cloud” thereby missing on important nuances about how discrete service models and delivery models can deliver different types of business value. Secondly, there were significant differences between IS and CS papers with regards to the granularity and substantiation of impact of the IT artefacts studied. While CS papers examined IT artefacts at an extremely low level of granularity in a cloud solution stack when compared to IS papers, they could clearly link these artefacts across the causal chain to economic factors in a way that IS papers could not or did not. Furthermore, the impact was measurable in much shorter time horizons. Thirdly, the techniques used to measure business value were concentrated on measuring costs e.g. Total Cost of Ownership (TCO). This is not wholly unsurprising given that the focus was mostly IaaS and migration from on-premise. However, more concerning was that many of the studies, and, in particular, CS studies, demonstrated significant methodological issues in their calculation of costs, and where examined, benefits. In particular, few attempts were identified to measure intangible benefits.

In summary, our feeling in 2017 was that there was a need for a more systematic and interdisciplinary approach to researching the conceptualisation and measurement of the business value of cloud computing (BVCC) in a more disaggregated way. Assuming an unchanging technological

landscape, this would have still been a major challenge. However, our conceptualisation of the “cloud” is radically changing. The pace of change in cloud computing and how enterprises manage and use it has accelerated dramatically in recent years. As a consequence, it is surely worth considering whether the nature of the business value created by cloud computing and how we measure it has changed too. This chapter presents a number of new paradigms in cloud computing, changes in cloud architectures, and research pathways we believe may prove promising avenues for future research for both IS and CS researchers.

7.2 THE CHANGING NATURE OF THE CLOUD

The accepted definition of cloud computing has not changed. It is:

...a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. (Mell and Grance 2011, p. 2)

However, the nature of the cloud has changed. It is increasing more abstracted, heterogeneous, composable, and automated.

7.2.1 *The Evolution of Shared Resources*

How resources are shared in the cloud is evolving rapidly (see Fig. 7.1). In the first phase of cloud computing, we saw a shift from monolithic architectures to service oriented architectures; this is what is largely described in the NIST Cloud Computing Reference Architecture (Mell and Grance 2011) and the focus of BVCC research from 2009 to 2016. In this phase, cloud service providers and their customers benefit from their own discrete virtual machines (VMs) running on shared infrastructure.




Fig. 7.1 Evolution of shared resources in cloud computing. Grey areas are shared. (Adapted from Hendrickson et al. 2016)

Since the open sourcing of dotCloud's container technology in 2013, the nature of the cloud began changing again. Containerisation enabled operating system (OS)-level virtualization where containers hold all the components necessary to run a specific software program and a minimal subset of an OS. As a concept, this results in a number of benefits relevant to measuring business value. For example, containers are less resource-intensive than VMs and therefore result in reduced operational expenditure. They are more portable thus reducing lock-in and increasing agility and flexibility. New services can be provisioned faster thus resulting in increased time to market. Despite these advantages, there is scant discussion of containerisation (or micro-services) in the IS literature and even less, if any, on the measurement of the business value of this architectural style.

More recently, we have seen the emergence of serverless cloud computing. Here, effectively all resources are pooled including hardware, operating systems and runtime environments. Serverless computing is "a software architecture where an application is decomposed into 'triggers' (events) and 'actions' (functions), and there is a platform that provides a seamless hosting and execution environment" (Glikson et al. 2017, p. 1). The software owner does not necessarily have to concern themselves with management of the runtime environment instead can focus on developing and deploying relatively lightweight, single purpose stateless functions that can be executed on-demand, typically through an API, without consuming any resources until the point of execution (Lynn et al. 2017). As such, this cloud service model is often called Function as a Service (FaaS). The cloud service provider assumes responsibility for data centre management, server management and the runtime environment. The software operators only pay for resources when they are executed thus reducing the cost of deployment dramatically. Furthermore, FaaS also transforms the business model of cloud service providers e.g. pricing at the level of execution runtime for computer code rather than how long an instance is running (Eivy 2017). For these reasons, FaaS is gaining significant traction. It has been adopted not only by the major hyperscale cloud service providers (e.g. Google, Microsoft, Amazon Web Services (AWS), and IBM) but also many well-known companies e.g. Netflix (transcoding, monitoring, disaster recovery, and compliance), Seattle Times (image resizing), Zillow (real-time mobile metrics), and Major League Baseball Advanced Media (data analysis, and player and game metrics) (Lynn et al. 2017). As previously mentioned, there was an existing need for more BVCC research relating to traditional

cloud computing service models i.e. IaaS, PaaS, and SaaS (and to a much lesser extent Business Process as a Service—BPaaS); there is virtually no research on measuring the business value of containerisation (microservices) or serverless cloud computing (FaaS).

7.2.2 *The Heterogeneous Cloud*

As cloud computing continues to become the dominant computing paradigm, cloud service providers are looking for new segments for growth and enterprises are looking for new ways to create business value from migrating to the cloud. Two segments which have garnered a lot of attention in recent years are Big Data analytics and high performance computing (HPC) in the cloud. The benefits of Big Data and related analytics include increased agility (Ashrafi et al. 2019), innovation (Lehrer et al. 2018), and competitive performance (Côte-Real et al. 2017; Mikalef et al. 2019) and is widely discussed in both IS and CS literature, and even more so in practice. The contribution of HPC is less widely discussed yet is recognised as playing a pivotal role in both science discovery and national competitiveness (Ezell and Atkinson 2016). The widespread use of both Big Data analytics and HPC have been hampered by significant upfront investment and indirect operational expenditure (including specialised staff) associated with running and maintaining these infrastructures. Big Data analytics and HPC in the cloud represent massive opportunities to unleash business value through reduced CapEx and OpEx as well as democratising Big Data and HPC infrastructure and tools and thus increase innovation output.

Traditionally, and to a large extent today, cloud computing systems are optimised to cater for multiple tenants and a large number of small workloads. The primary focus of traditional cloud computing is rapid scalability and as such is designed for perfectly or pleasingly parallel problems (Lynn 2018). For such workloads, while servers must be available and operational, neither the precise physical server nor the speed of the connections between processors that executes a request is important provided the resource database remains coherent (Eijkhout et al. 2016). Unfortunately, for Big Data analytics or HPC workloads, enterprise users typically require servers to be available on-demand and connected via high-speed, high-throughput, and low-latency network interconnects (Lynn 2018).

Heterogeneous computing refers to architectures that allow the use of different hardware types to work efficiently and cooperatively together.

Unlike traditional cloud infrastructure built on the same processor architecture, heterogeneity assumes use of different or dissimilar processors or cores that incorporate specialised processing capabilities to handle specific tasks more faster and more energy efficiently than general purpose processors (Scogland et al. 2014). For example, field-programmable gate arrays (FPGAs) and graphics processing units (GPUs) are co-processor architectures with relatively positive computation/power consumption ratios that offer significant performance and energy efficiency gains for Big Data analytics and HPC respectively. Increasingly heterogeneous computing is being extended beyond different processor architectures to include different networking infrastructure that can support higher throughput and lower latency (Shan 2006; Yeo and Lee 2011). In recent years, major public cloud providers including AWS, Microsoft Azure and Google Cloud offer specialist cloud services for Big Data and HPC uses cases built on heterogeneous clouds. These specialist clouds are increasingly being adopted by some of the world's largest companies including Aon (financial simulation), AstraZeneca (genome processing), BP (linear programming models), Disney (video streaming analytics and rendering), and Volkswagen (computation fluid dynamics). Despite the increasing availability and use of heterogeneous cloud computing, there is little research on the business value of adopting heterogeneous cloud computing.

7.2.3 *The Composable Cloud*

As more and more enterprises embrace digital transformation, even when private clouds are adopted, traditional IT architectures struggle to accommodate the cloud computing requirements from next generation applications. Legacy applications require infrastructure resiliency and exploit virtualization and clustering for portability and application state preservation. In contrast, next generation applications (NGAs) are designed to be horizontally scalable, containerised and continuously updated (Nadkarni 2017). IDC suggest that in most enterprise data centres, infrastructure is 45% over provisioned, 45% utilized, and 40% compliant with stated service level agreements (Nadkarni 2017).

Composable architectures assume that resources (e.g. compute, memory, storage, networking etc.) can be decoupled from the hardware they reside on and assembled and re-assembled using a control software layer to meet exact workload requirements on-demand (Ferreira et al. 2019). Once hardware is no longer required it can be released for use for another

workload. There are a number of advantages to this approach. Firstly, discrete servers do not need to be configured for a specific application but rather hardware resources can be pooled to meet both legacy and NGAs dynamically. If more resources are needed to deliver a given workload, it is automatically provisioned. Secondly, composable architectures support heterogeneous computing and pools these resources in the same way thereby allowing enterprises to exploit the performance or energy efficiencies of these specialist resources. Thirdly, as each workload is provisioned exactly as needed, over-provisioning is reduced dramatically hereby reducing both CapEx and OpEx.

The Composable Cloud is a fundamentally different way to operate data centres and private clouds. Given that it reduces overprovisioning and related inefficiencies dramatically as well as freeing up valuable enterprise resources, not least cash flow and staffing, it is worthy of investigation by business value researchers.

7.2.4 *The Automated Cloud*

A side effect of new service models, increased heterogeneity, and composable is greater complexity in terms of reliability, maintenance and security (Marinescu 2017). This is particularly the case for large-scale enterprise systems and hyperscale cloud services where the scale of infrastructure, applications, and number of end users is significant. It is no longer feasible for IT teams to cost-effectively foresee and manage manually all possible configurations, component interactions, and end-user operations on a detailed level due to high levels of dynamism in the system (Lynn 2018). As such, enterprise IT and cloud service providers are increasingly looking to machine learning and artificial intelligence (AI) to manage this complexity but also automate previously manual tasks, and free up staff.

AIOps (AI for IT Operations) uses algorithms and machine learning to dramatically improve the monitoring, operation, and maintenance of distributed systems (Cordoso 2019). The main use cases for AIOps are performance analysis, anomaly detection, event correlation and analysis, and IT service management and automation with the ultimate goals of ensuring high service quality and customer satisfaction, boosting engineering productivity, and reducing operational costs (Prasad and Rich 2018; Dang et al. 2019). IDC's Worldwide Developer and DevOps 2019 Predictions suggest that by 2024, 60% of firms will have adopted AIOps (Gillen et al. 2018). Much of the market demand for AIOps is couched in

the fear of outages and the ability of machine learning to predict such outages and enable preventative action to be taken before customers or business is impacted. Yet despite this optimism, there are significant challenges with the adoption of AIOps including changes in innovation methodologies including understanding business value and constraints, engineering mindsets, and engineering practices including data quality (Dang et al. 2019). From a business value research perspective, machine learning and AI poses additional challenges as the black box nature of these technologies can make interrogation and interpretability difficult.

7.3 CLOUD COMPUTING AND THE INTERNET OF THINGS

Over the last five years, interest in the Internet of Things (IoT) has increased dramatically, partially fuelled by the increasing ubiquity of internet access and smartphones but also estimates of the value of IoT forecast to exceed \$19 trillion over time (Cisco 2013a, b). This value is generated through connecting a fraction of the 1.4 trillion things *in situ* today, and consequently improving asset utilization, employee productivity, supply chain and logistics, customer experience, as well as accelerating innovation (Lynn et al. 2018).

Haller et al. (2009) define IoT as:

A world where physical objects are seamlessly integrated into the information network, and where the physical objects can become active participants in business processes. Services are available to interact with these “smart objects” over the Internet, query their state and any information associated with them, taking into account security and privacy issues. (Haller et al. 2009, p. 15)

Smart objects may range from sensors, with little storage and data processing power, to modern smartphones. IoT assumes that smart things can carry out, with minimal latency, some degree of data processing and collaborate with other devices and systems, some local and some remotely. As such, it assumes a continuum of computing activity from the cloud to thing (C2T) where computing resources can be located in the cloud, at the thing (edge computing), or somewhere in between (fog computing). As such, IoT effectively extends cloud computing from a centralised service architecture to a decentralised one. Table 7.1 below summarises key definitions of new computing paradigms along the C2T continuum.

Table 7.1 Definitions of edge, fog and mist computing. (Adapted from Iorga et al. 2018)

<i>Concept</i>	<i>Definition</i>
Edge computing	Edge computing is the network layer encompassing the end devices and their users, to provide, for example, local computing capability on a sensor, metering or some other devices that are network-accessible.
Fog computing	Fog computing is a layered model for enabling ubiquitous access to a shared continuum of scalable computing resources. The model facilitates the deployment of distributed, latency-aware applications and services, and consists of fog nodes (physical or virtual), residing between smart end-devices and centralized (cloud) services.
Mist computing	Mist computing is a lightweight and rudimentary form of fog computing that resides directly within the network fabric at the edge of the network fabric, bringing the fog computing layer closer to the smart end-devices. Mist computing uses microcomputers and microcontrollers to feed into fog computing nodes and potentially onward towards the centralized (cloud) computing services.

For enterprises, cloud service providers and cloud carriers (e.g. Tier 1 network operators), IoT introduces complexity at yet another higher order of magnitude. To meet the Quality of Service (QoS) and Quality of Experience (QoE) requirements of SLAs with customers and/or end users, service providers and cloud carriers need to decide where best to locate compute and storage resources along the C2T continuum. As such, enterprises need to consider the geographic distribution and mobility of smart objects and latency at each location, the heterogeneity of smart objects, interoperability and federation, the necessity and capability of real-time interaction, and the scalability and agility of federated fog-node clusters (Iorga et al. 2018).

Haller et al. (2009) suggest that there are two main sources where enterprises can derive business value from the IoT—real world visibility and business process decomposition. Firstly, they argue that the use of automated identification and data collection technologies will give enterprises unparalleled insights in to what is happening in the real world thus enabling high resolution management and the potential deeper and better business insights, more effective optimisation, and better decision making. Secondly, they argue that IoT combined with real world visibility allows the decomposition of business processes in to process steps (and associated computing resources) which can be distributed from the cloud to edge

thus enabling the decentralisation of business processes resulting in increased scalability and performance, better decision making, and innovation. From a cloud computing perspective, IoT involves key technical decisions that can impact the business value generated for the enterprise e.g. how much infrastructure should be placed at different points across the C2T continuum? What applications (or if distributed, what application components) should be operated at the edge and which should not? How do these placement decisions impact business value?

7.4 TOWARDS AN AGENDA FOR BUSINESS VALUE IN CLOUD COMPUTING RESEARCH

Cloud computing is a key enabling technology. As can be seen above, its interaction with other technologies, for example machine learning and AI, mobile computing, IoT, and HPC accelerates innovation and as a result potential business value. Recently, a number of authors have suggested research gaps and questions that can guide future research on business value in cloud computing. In particular, Schryen (2013) calls for greater research to close three IS business value research gaps which are relevant to business value in cloud computing i.e. ambiguity and fuzziness of the ‘IS business value’ construct, neglected disaggregation of IS investments, and the IS business value creation process as a grey box. Paraphrasing Schryen, a number of research pathways for business value in cloud computing arise:

How can we yield a comprehensive, consistent and precise understanding of the multifaceted construct ‘cloud computing business value’? How can the assessment of (internal and competitive) business value account for the context of evaluation, and in particular the firm, industry, and country environment and preferences of evaluators? (Schryen 2013, pp. 151–152)

Schryen (2013) suggests ‘IS business value’ is ambiguous and fuzzy. As previously discussed, our own experience is that not only is IS business value ambiguous and fuzzy but the techniques for measuring business value are often ambiguous and, where documented, are not applied consistently or comprehensively in such a way to allow comparison. In addition, a more nuanced approach to defining the firm and industry is needed. At a basic level, firms may include enterprises adopting cloud computing, cloud service providers, cloud carriers (e.g. network operators), cloud brokers, cloud auditors, and indeed edge device

consumers (See Fig. 7.2 below). Each of these actors may operate in different industries and thus provide different industry and country contexts and associated constraints, particularly with respect to operation. For example, cloud computing is, by and large, a cross-border phenomenon however national data privacy laws, amongst others, create opportunities and risks for business value generation and capture.

How can total cloud computing investments be disaggregated conceptually and empirically such that the impact of different types of investments on the economic performance of the firm can be determined? How can the disaggregation of total cloud computing investments account for synergies and complementarities? (adapted from Schryen 2013, pp. 153–154)

This assumes one can disaggregate cloud computing investments from wider IS investments and then specific cloud computing investments. As indicated earlier, at a basic level this could be by service model (IaaS, PaaS, SaaS, and FaaS) or even by deployment model (private, public, hybrid, and community clouds), components of the extended cloud computing conceptual reference model (see Fig. 7.2), or a combination of all of these.




Fig. 7.2 Extended cloud computing conceptual reference model. (Adapted and extended from Liu et al. 2011)

To address these research questions, it assumes (1) a sufficiently detailed taxonomy of cloud computing investments can be catalogued, (2) critical success factors (CSFs) and key performance indicators (KPIs) can be mapped to these supporting assets, and (3) occurrences of synergies between different types of assets can be identified (Schryen 2013). This may require examination at a lower level of granularity than IS researchers typically undertake and as such may require CS support thus mandating interdisciplinary research.

How, why and when do cloud computing assets, cloud computing capabilities, IS assets and capabilities, and socio-organisational capabilities affect each other and joint create internal value? How, why, and when do cloud computing assets, cloud computing capabilities, IS assets and capabilities, and socio-organisational capabilities create competitive value, thus performing a value creation process? (adapted from Schryen 2013, p. 156)

This research question recognises that cloud computing assets and capabilities are a subset of wider IS assets and capabilities and have a bidirectional relationship with socio-organisational capabilities. This is particularly the case when we consider emerging use cases including IoT, Big Data analytics and HPC. It also recognises that value is created over time and that some aspects are static and some are dynamic. In the context of cloud computing, firms more than likely inherit the assets and capabilities of the chain of service provision and thus for a given time, have compound capabilities or what Carroll et al. (2013, 2014) call a composite capability. The business value of such capabilities is dependent on a number of socio-organisational factors, not least size, which obviously changes over time. As such, research must consider a contingency approach to business value.

7.5 CONCLUDING REMARKS

This chapter presents a number of new paradigms in cloud computing, changes in cloud architectures, and research pathways in business value in cloud computing research that we believe may provide future avenues of research for both IS and CS researchers. This is by no means exhaustive. Indeed, other chapters in this book cover aspects of business value in cloud computing research that could provide a fruitful stream of research. As we develop our understanding of cloud computing and the dependencies between cloud computing and other technologies (not least mobile, Big

Data, and IoT) the need for greater clarity on the definition and appropriate metrics of business value; robust business value measurement techniques; disaggregation of IS assets (and specifically cloud infrastructure); and the relationship between cloud assets and capabilities, other IS assets and capabilities, and socio-organisation capabilities, is required. This will require a deep understanding of these technologies and most likely collaboration between information systems and computer science researchers. More importantly, it will require a change in the mindsets of business value researchers in both disciplines.

REFERENCES

- Ashrafi, A., A.Z. Ravasan, P. Trkman, and S. Afshari. 2019. The Role of Business Analytics Capabilities in Bolstering Firms' Agility and Performance. *International Journal of Information Management* 47: 1–15.
- Carroll, N., M. Helfert, and T. Lynn. 2013. *A Contingency Model for Assessing Cloud Composite Capabilities*. CLOSER, 515–519.
- . 2014. Towards the Development of a Cloud Service Capability Assessment Framework. In *Continued Rise of the Cloud*, 289–336. London: Springer.
- Cisco. 2013a. Internet of Everything: A \$4.6 Trillion Public-Sector Opportunity. https://www.cisco.com/c/dam/en_us/about/business-insights/docs/ioe-public-sector-vas-white-paper.pdf.
- . 2013b. Embracing the Internet of Everything to Capture Your Share of \$14.4 Trillion. https://www.cisco.com/c/dam/en_us/about/business-insights/docs/ioe-economy-insights.pdf
- Cordoso, J. 2019. *The Application of Deep Learning to Intelligent Cloud Operation*. Paper presented at Huawei Planet-scale Intelligent Cloud Operations Summit, Dublin, Ireland, 1 November 2019.
- Côrte-Real, N., T. Oliveira, and P. Ruivo. 2017. Assessing Business Value of Big Data Analytics in European Firms. *Journal of Business Research* 70: 379–390.
- Dang, Y., Q. Lin, and P. Huang. 2019. *AIOPs: Real-World Challenges and Research Innovations*. Proceedings of the 41st International Conference on Software Engineering: Companion Proceedings, 4–5. IEEE Press.
- Eijkhout, V., R. van de Geijn, and E. Chow. 2016. *Introduction to High Performance Scientific Computing*. Zenodo. <https://doi.org/10.5281/zenodo.49897>
- Eivy, A. 2017. Be Wary of the Economics of “Serverless” Cloud Computing. *IEEE Cloud Computing* 4 (2): 6–12.
- Ezell, S.J., and R.D. Atkinson. 2016. *The Vital Importance of High Performance Computing to US Competitiveness*. Washington, DC: Information Technology

- and Innovation Foundation. Accessed 23 October 2017. <http://www2.itif.org/2016-high-performance-computing.pdf>
- Ferreira, L., E. da Silva Rochay, K. Monteiro, G. Santos, F. Silva, J. Kelner, D. Sadok, C. Bastos Filho, P. Rosati, T. Lynn, and P. Endo. 2019. *Optimizing Resource Availability in Composable Data Center Infrastructures*. Proceedings of the 9th Latin-American Symposium on Dependable Computing. IEEE.
- Gillen, A., C. Arend, M. Ballou, L. Carvalho, A. Dayaratna, S. Elliot, M. Fleming, M. Iriya, P. Marston, J. Mercer, G. Mironescu, J. Thomson, and C. Zhang. 2018. *IDC FutureScape: Worldwide Developer and DevOps 2019 Predictions*. IDC.
- Glikson, A., S. Nastic, and S. Dustdar. 2017. *Deviceless Edge Computing: Extending Serverless Computing to the Edge of the Network*. Proceedings of the 10th ACM International Systems and Storage Conference, 28. ACM.
- Haller, S., S. Karnoukos, and C. Schroth. 2009. *The Internet of Things in an Enterprise Context*. Future Internet Symposium, 14–28. Berlin, Heidelberg: Springer.
- Hendrickson, S., Sturdevant, S., Harter, T., Venkataramani, V., Arpaci-Dusseau, A. C. and Arpaci-Dusseau, R. H. (2016). Serverless Computation with OpenLambda. *Elastic* 60: 80.
- Iorga, M., L. Feldman, R. Barton, M. J. Martin, N. S. Goren, and C. Mahmoudi. 2018. Fog Computing Conceptual Model (No. Special Publication (NIST SP)-500-325).
- Lehrer, C., A. Wieneke, J. vom Brocke, R. Jung, and S. Seidel. 2018. How Big Data Analytics Enables Service Innovation: Materiality, Affordance, and the Individualization of Service. *Journal of Management Information Systems* 35 (2): 424–460.
- Liu, F., J. Tong, J. Mao, R. Bohn, J. Messina, L. Badger, and D. Leaf. 2011. NIST Cloud Computing Reference Architecture: Recommendations of the National Institute of Standards and Technology (Special Publication 500–292).
- Lynn, T. 2018. Addressing the Complexity of HPC in the Cloud: Emergence, Self-organisation, Self-management, and the Separation of Concerns. In *Heterogeneity, High Performance Computing, Self-Organization and the Cloud*, 1–30. Cham: Palgrave Macmillan.
- Lynn, T., P. Rosati, A. Lejeune, and V. Emeakaroha. 2017. *A Preliminary Review of Enterprise Serverless Cloud Computing (Function-as-a-Service) Platforms*. 2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom), 162–169. IEEE.
- Lynn, T., P. Rosati, and P. Endo. 2018. *Towards the Intelligent Internet of Everything: Observations on Multi-disciplinary Challenges in Intelligent Systems*. Proceedings of the Research Coloquio Doctorados: Tecnología, Ciencia y Cultura: una visión global.
- Marinescu, D.C. 2017. *Complex Systems and Clouds: A Self-organization and Self-management Perspective*. Cambridge, MA: Elsevier.

- Mell, P., and T. Grance. 2011. *The NIST Definition of Cloud Computing*, Special Publication 800-145. Gaithersburg, MD: National Institute of Standards and Technology.
- Mikalef, P., J. Krogstie, I.O. Pappas, and P. Pavlou. 2019. Exploring the Relationship between Big Data Analytics Capability and Competitive Performance: The Mediating Roles of Dynamic and Operational Capabilities. *Information & Management* 57: 103169.
- Nadkarni, A. 2017. *Quantifying Datacenter Inefficiency: Making the Case for Composable Infrastructure*. IDC.
- Prasad, P., and C. Rich. 2018. *Market Guide for AIOps Platforms*. Gartner.
- Rosati, P., G. Fox, D. Kenny, and T. Lynn. 2017. *Quantifying the Financial Value of Cloud Investments: A Systematic Literature Review*. 2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom), 194–201. IEEE.
- Schryen, G. 2013. Revisiting IS Business Value Research: What We Already Know, What We Still Need to Know, and How We can Get There. *European Journal of Information Systems* 22 (2): 139–169.
- Scogland, T.R., C.P. Steffen, T. Wilde, F. Parent, S. Coghlan, N. Bates, et al. 2014. *A Power-Measurement Methodology for Large-Scale, High-Performance Computing*. Proceedings of the 5th ACM/SPEC International Conference on Performance Engineering, 149–159. ACM.
- Shan, A. 2006. Heterogeneous Processing: A Strategy for Augmenting Moore's Law. *Linux Journal* 2006 (142): 7.
- Yeo, S., and H.H. Lee. 2011. Using Mathematical Modeling in Provisioning a Heterogeneous Cloud Computing Environment. *Computer* 44 (8): 55–62.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

