

Scaling Edge Computing in the Real World

Key Considerations for Management, Security and Total Cost of Ownership



**Special
Report**

September 2021

Commissioned by

ZEDEDA

451 Research

S&P Global
Market Intelligence

©Copyright 2021 S&P Global Market Intelligence. All Rights Reserved.

About the Author



Christian Renaud

Global Head, 451 Advisors

Christian Renaud is the Global Head of 451 Advisors within S&P Global Market Intelligence, leading a team of consultants delivering bespoke engagements for research clients. He previously served as Research Director for the Internet of Things channel of 451 Research.

In the decades prior to joining S&P/451 Research, Christian built nationwide networks at large and small enterprises, worked with Fortune 50 companies in the systems integrator channel, built products at Cisco Systems and ran the company's New Markets and Technologies team. He has been the CEO of multiple startups, worked in venture and angel capital, and has served as an advisor to G20 and European Commission projects.

Executive Summary

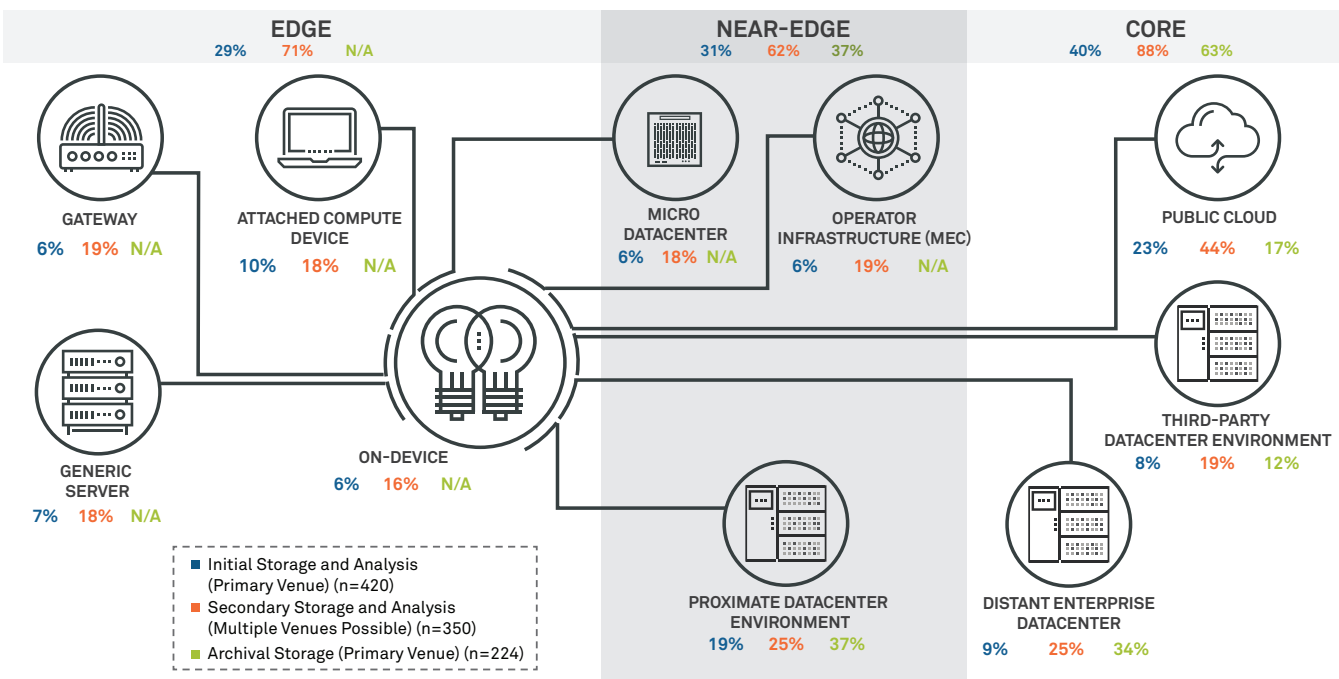
Edge computing is a powerful new tool in deploying mission-critical workloads. The combination of increasingly capable hardware with flexible, modular software built on a cloud-native foundation rounds out the continuum of compute and analytics from the point of data origination to the cloud. This paper discusses the practical realities of real-world edge implementations when moving beyond small proof-of-concept trials to large-scale deployments.

- The best execution venue for workloads in the continuum between the edge and the cloud is determined based on each organization's priorities for performance, cost and security.
- Use cases that leverage edge computing are frequently less expensive on a five-year basis when compared to cloud-centric models.
- Investing in open, standards-based edge management and security infrastructure prevents vendor lock-in and enables any permutation of edge or cloud compute models.

The Where, Why and ‘How Much?’ of the Edge

Edge deployments are already well underway, as evidenced by recent 451 Research survey data: 60% of respondents indicated that they are utilizing non-centralized computing locations for initial data analytics (see Figure 1). These locations are as diverse as embedded edge compute within larger machinery, stand-alone devices such as servers and gateways, and edge nodes within datacenters, carrier networks or content delivery networks.

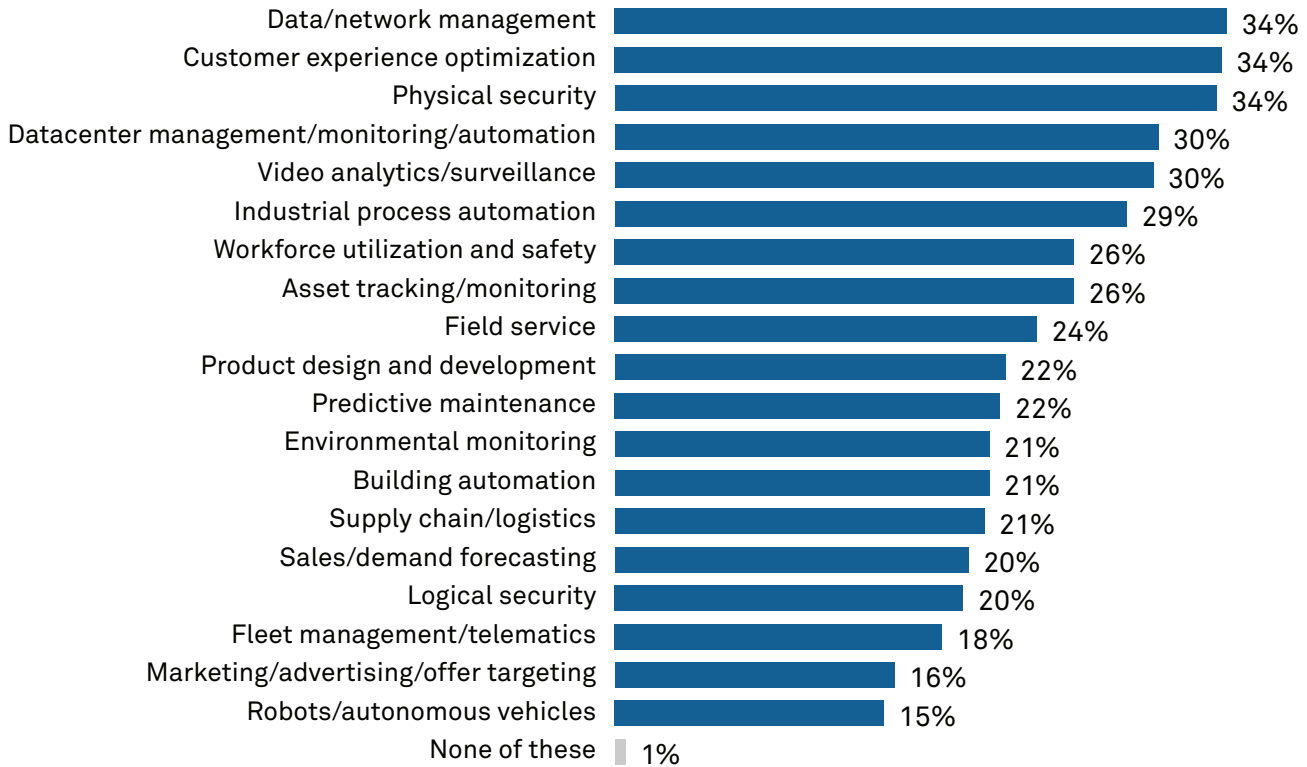
Figure 1: Initial and Secondary Storage and Analytics Venue



Source: 451 Research's Voice of the Enterprise: IoT, Workloads & Key Projects 2020

The applications at these execution venues vary by industry. However, a number of use cases, such as video analytics and data management, span multiple vertical markets (Figure 2) and generate considerable data, require low latency or have to adhere to strict data sovereignty regulations. The key drivers for edge deployments also vary by use case and vertical. They include the cost and availability of sufficient bandwidth, the latency requirements of real-time applications, security, data sovereignty, reliability for mission-critical workloads, and often being the physical bridge to legacy systems that are not IP/Ethernet-native. This is especially true in critical manufacturing, energy/oil and gas, and healthcare environments.

Figure 2: Use Cases by Industry and Affinity for Edge Compute

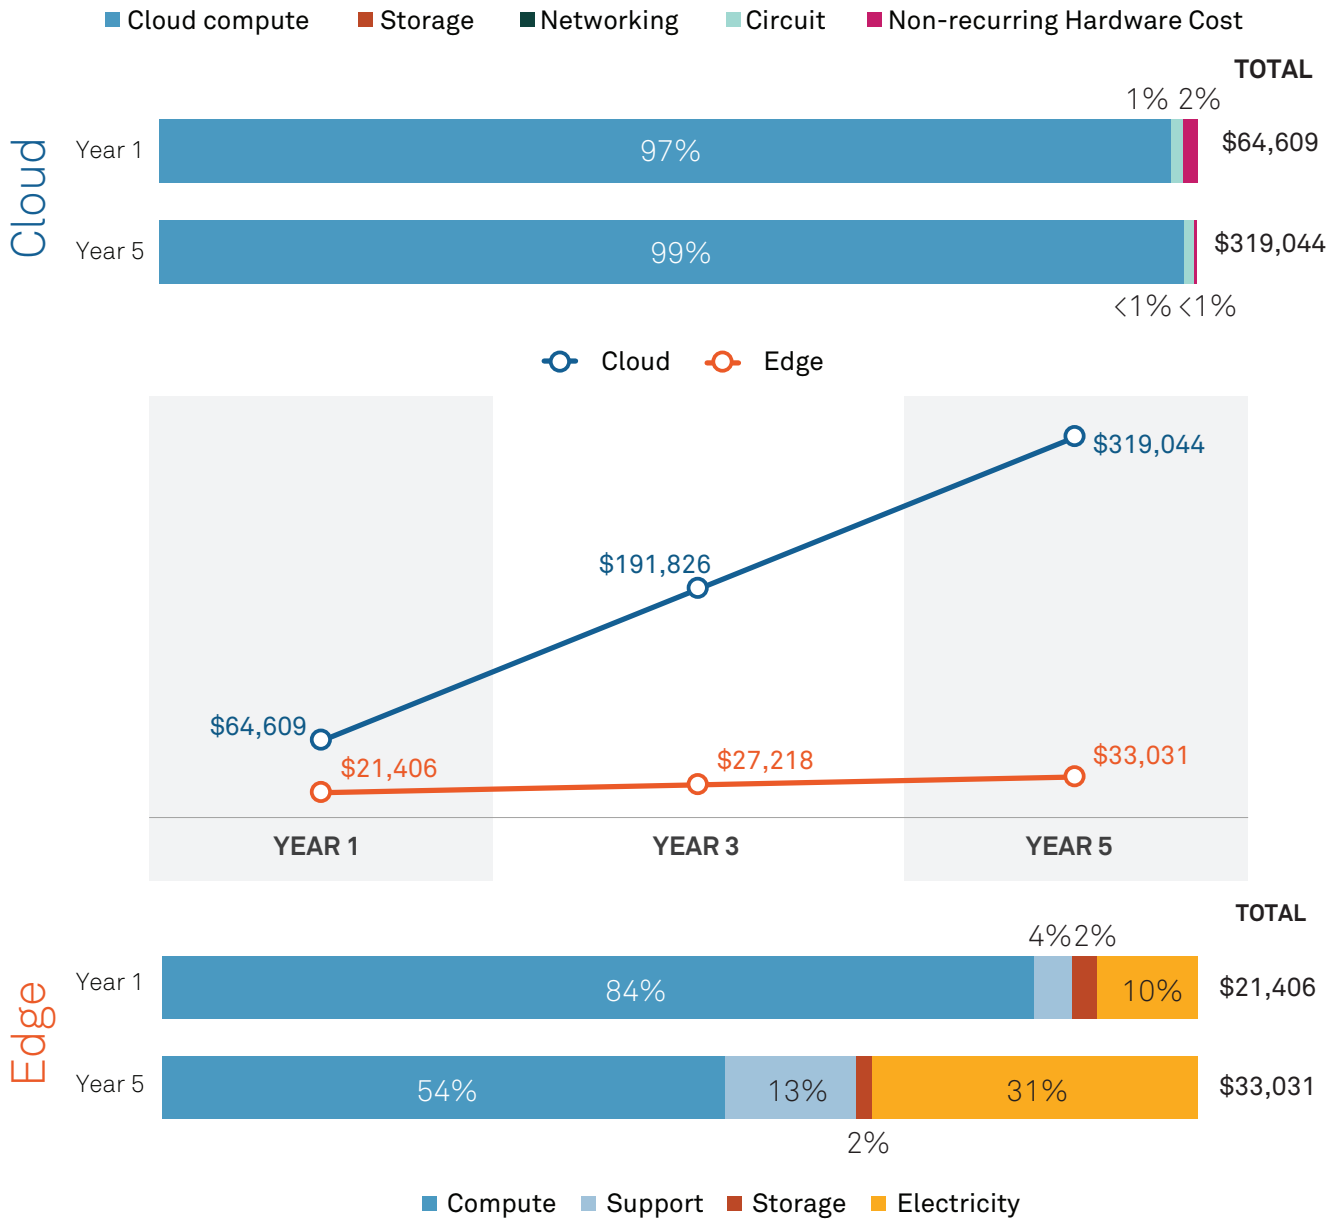


Source: 451 Research's Voice of the Enterprise: IoT, Workloads & Key Projects 2021

The edge is one end of a continuum of execution venues from the originating device generating data to a central cloud. This continuum gives practitioners considerable flexibility to place workloads where they can achieve the goals of performance, cost and security. Custom research conducted by 451 Research across multiple use cases in oil and gas, manufacturing and smart cities identified multiple key drivers that impact the total capital and operational expenditures over a five-year period. One use case from each industry is illustrated in Figure 3, comparing the total cost of leveraging the cloud or the edge for each workload. In each instance, the edge was the less expensive option over a five-year period, with the cost breakdown varying by workload depending on the volume of data, data retention (and thus storage required) policies, and circuit availability and cost.

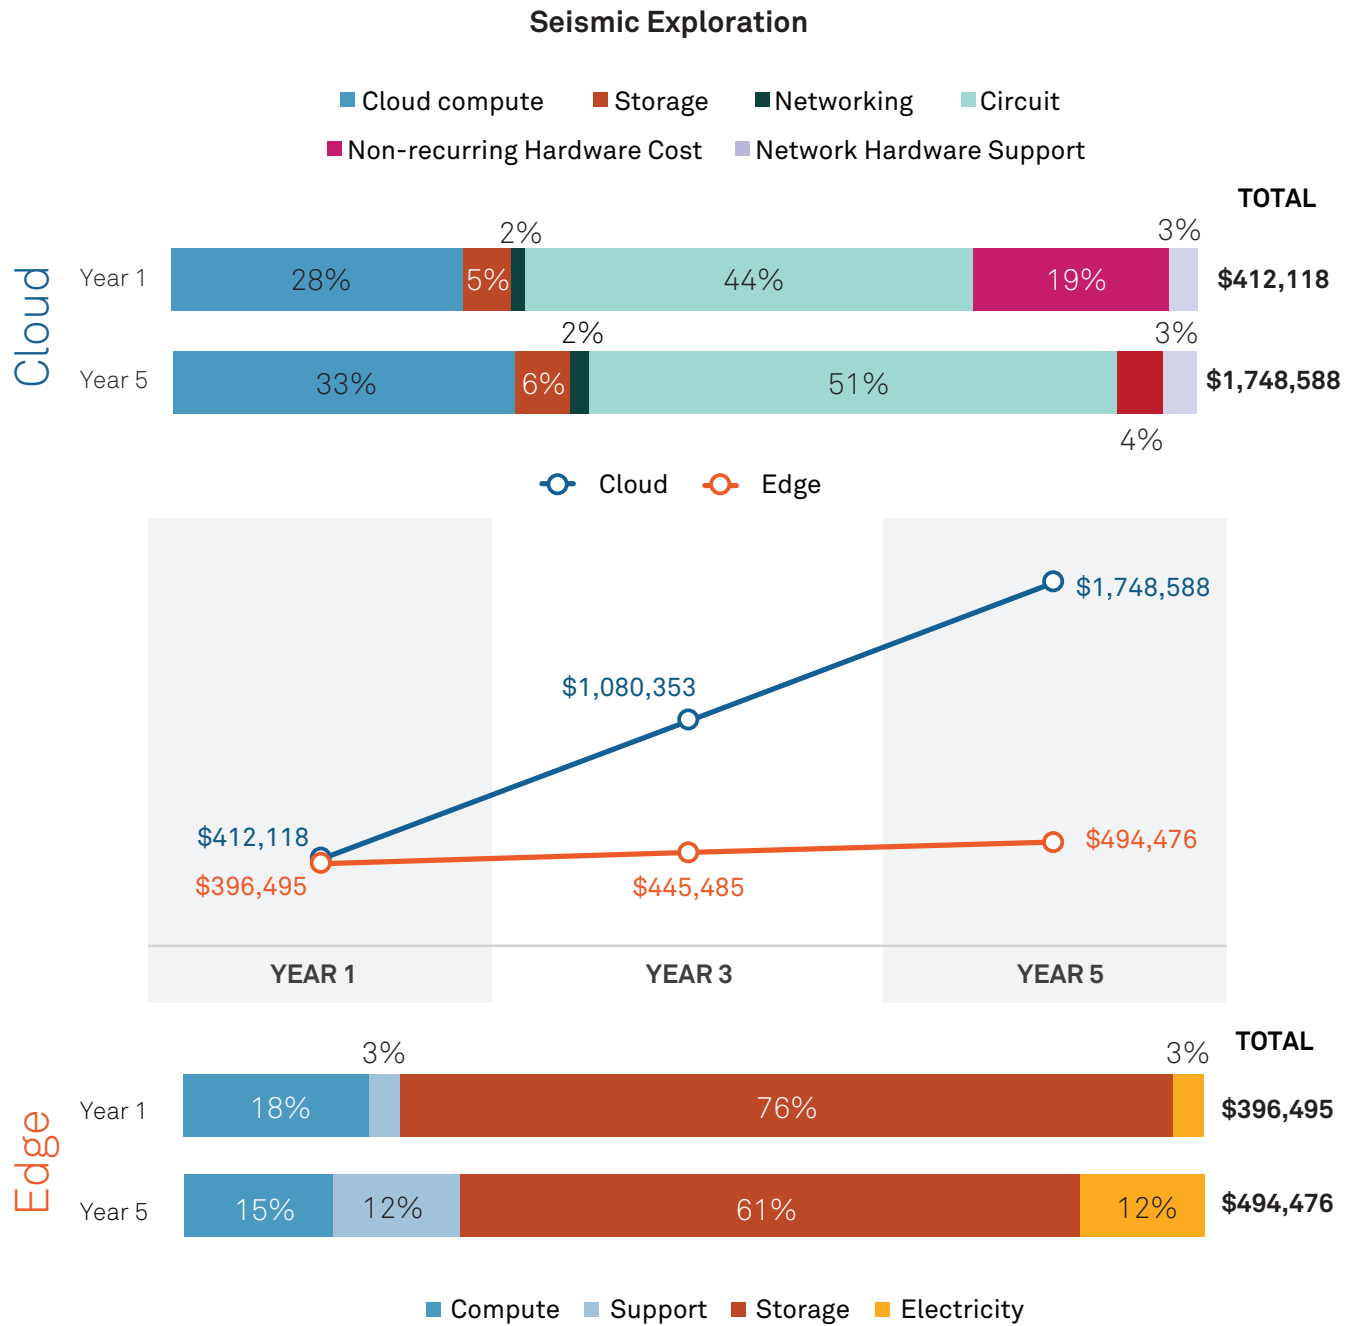
Figure 3-1: Five-Year Cost Breakdowns of Three Use Cases Comparing Cloud and Edge Execution Venues

Manufacturing Quality/Yield Monitoring



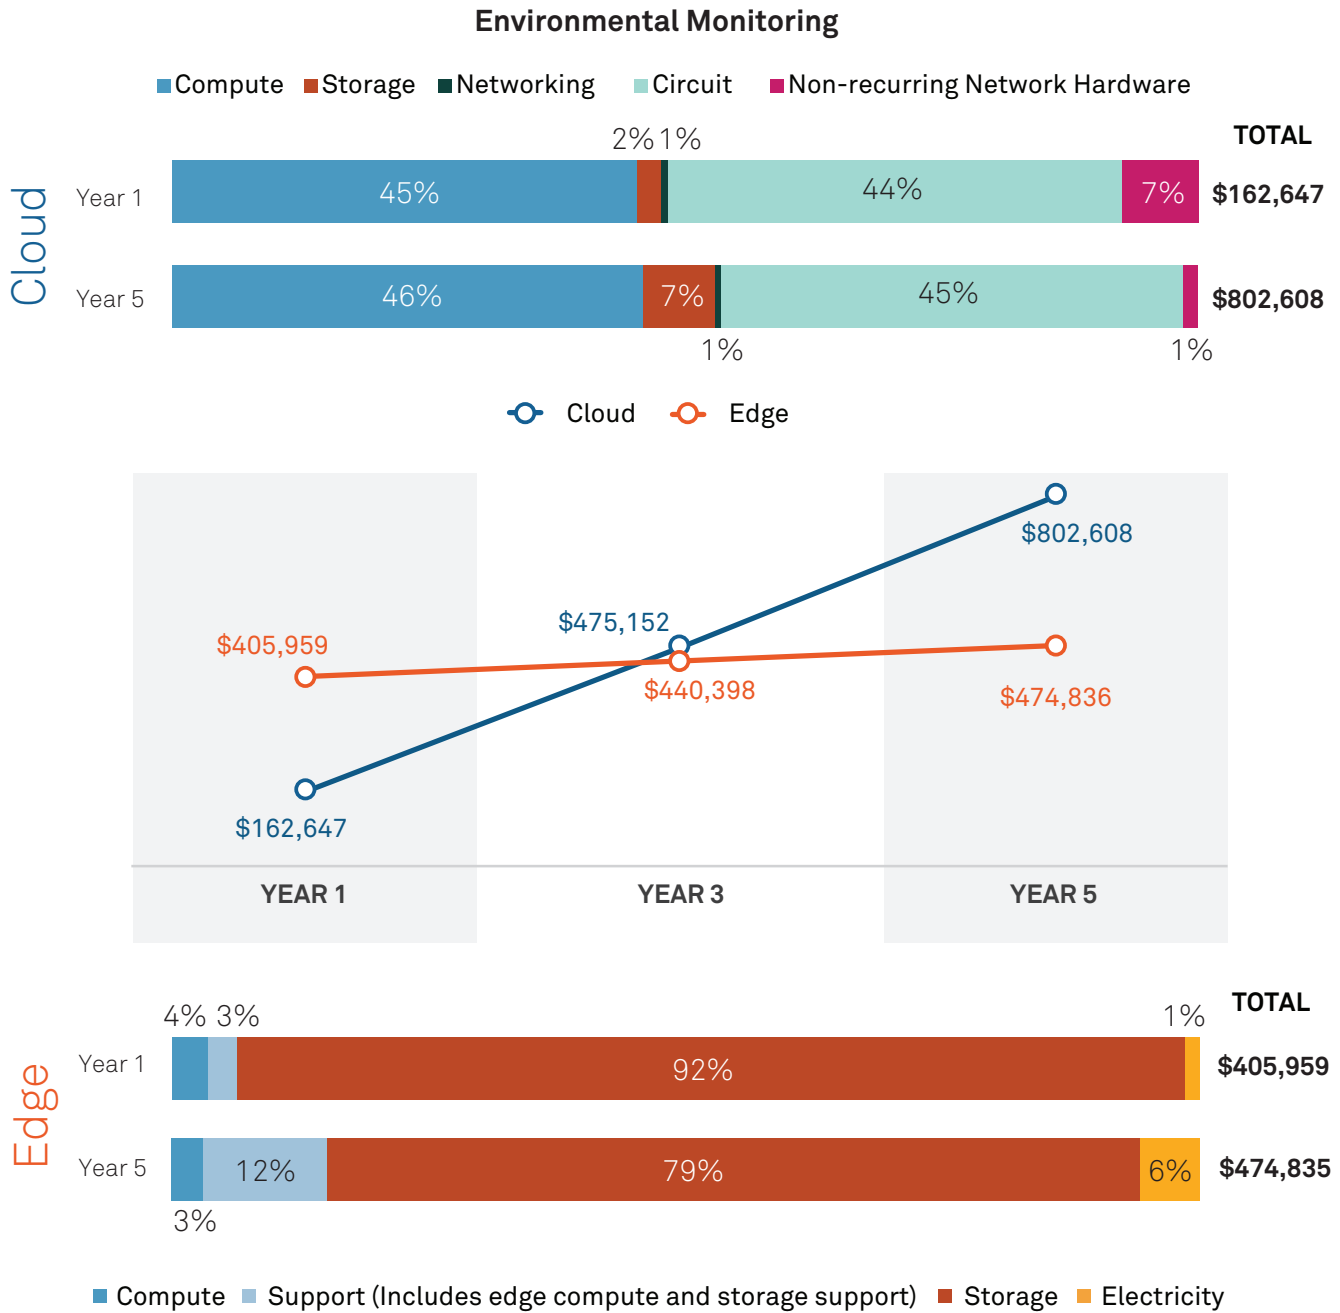
Source: 451 Research custom research

Figure 3-2: Five-Year Cost Breakdowns of Three Use Cases Comparing Cloud and Edge Execution Venues



Source: 451 Research custom research

Figure 3-3: Five-Year Cost Breakdowns of Three Use Cases Comparing Cloud and Edge Execution Venues

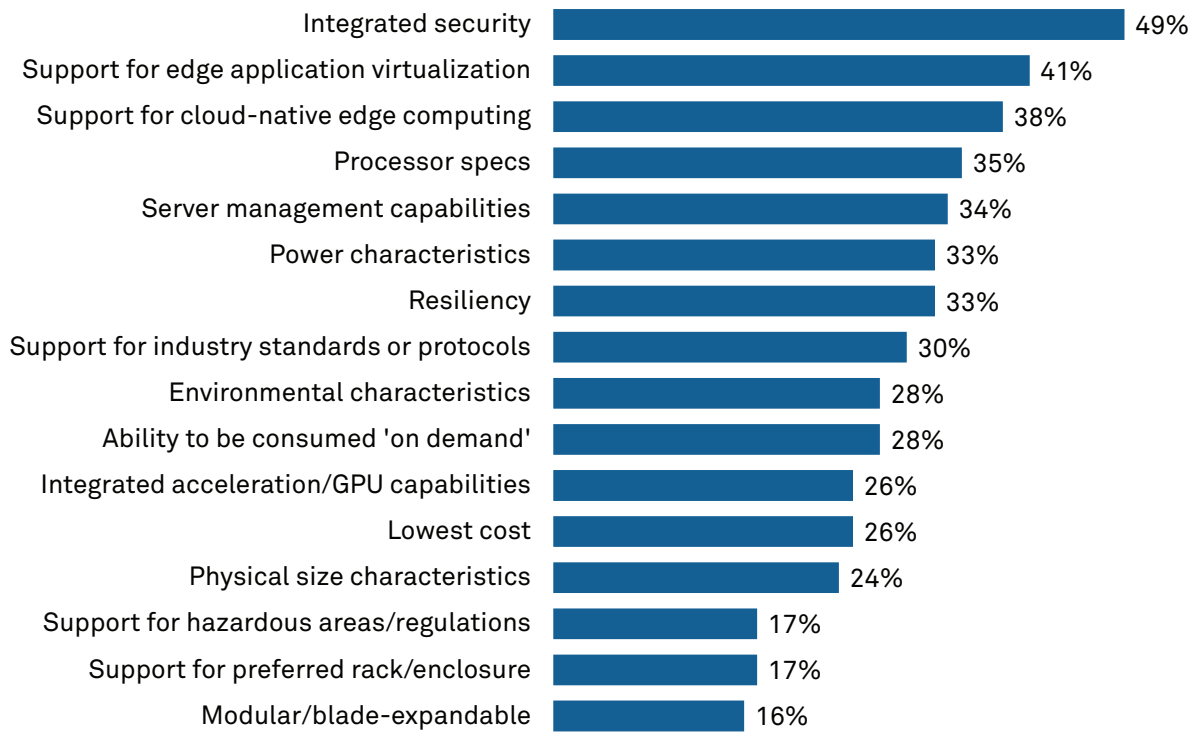


Source: 451 Research custom research

Executing a Successful Edge Deployment

Just as the cloud introduced new disciplines and phrases into the IT lexicon, so does the edge. Early edge implementations on compute-constrained hardware and proprietary software exposed limitations that have been addressed as the industry evolved from early proof-of-concept to large-scale deployments. Enterprises have become more discerning when selecting edge technologies and now prioritize security and built-in flexibility to avoid costly truck rolls and downtime.

Figure 4: Most Important Edge Product Attributes



Source: 451 Research's Voice of the Enterprise: IoT, Budgets and Outlook 2021

The costs of deploying, maintaining and securing the edge can far outweigh the costs of the initial hardware and software. The first of these chronologically is initial deployment and provisioning. As enterprises scale from tens of devices to hundreds or thousands of edge devices, the process of manually configuring them quickly 'fails to scale' and also introduces the potential for security lapses. Having a robust orchestration system that can provision an entire fleet of edge devices, as well as enforce a consistent security posture, is mandatory to avoid runaway operational expenses in deploying large quantities of geographically distributed edge devices.

Second is manageability – whether the broader deployment can be orchestrated as a single system. This has historically broken down due to proprietary silos and lack of data exchange among vendors, a catalyst for the creation of numerous industry efforts to break down the walls to multi-vendor interoperability. This also mandates a mature management scheme that discriminates and separates data traffic from management signaling. The combination of open, interoperable orchestration and the bifurcation of the application and orchestration planes builds a robust foundation for large-scale edge deployments that maximizes uptime and remote administration while minimizing physical touch points and truck rolls.

The third – albeit equally important – cost is ensuring that the systems deployed are physically and logically secure. Edge devices are frequently deployed outside of traditional physically secure and climate-controlled datacenters in production environments, which makes them attractive targets for security hacks, thus requiring a zero trust security model to address. At the same time, legacy OT infrastructure frequently predates best IT practices for logical security and segmentation. Best-of-class edge deployments have both operating system and application update and patch processes and protocols, and they also perform ongoing monitoring of edge devices and software for anomalous behavior that ventures outside of established baselines.

Looking Forward

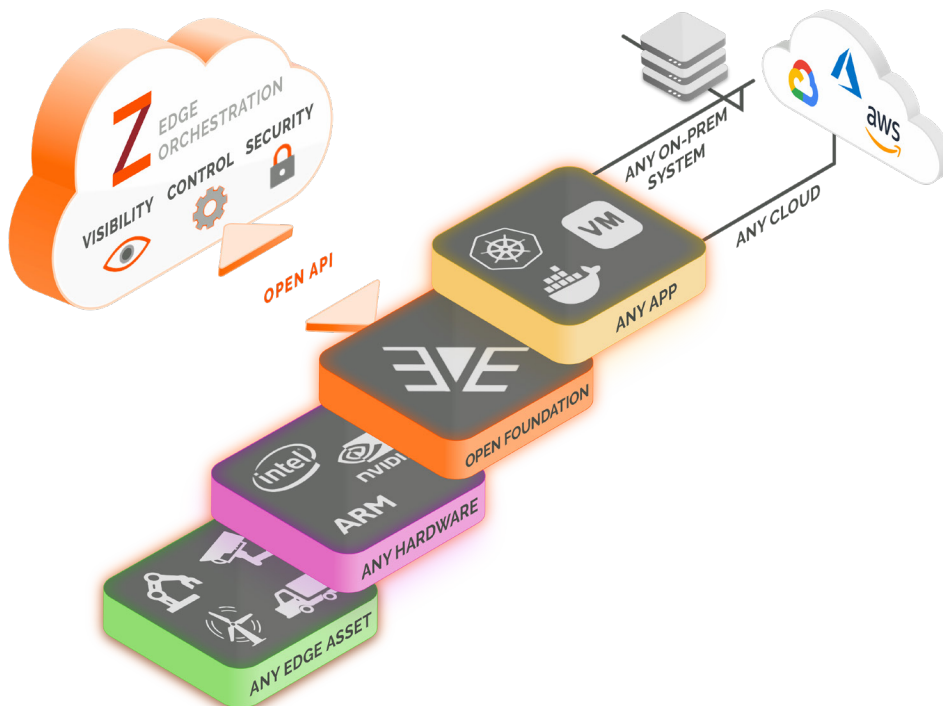
Edge technologies have quickly matured from hobbyist to mainstream adoption, but not all edge platforms address the entire spectrum of needs – provisioning, security and ongoing lifecycle management. Some offerings remain highly proprietary, a traditional hurdle in the maturity of any technology, and one that can easily lead enterprises into a dead-end road of lack of interoperability between vendors. Organizations including the Linux Foundation and Eclipse Foundation have developed open edge infrastructure implementations specifically to avoid this.

Edge technologies hold the promise of cost-effective and performant analytics close to the point of data origination, enabling a DataOps culture of putting the right data in the hands of the right people at the right time. The transition from early edge implementations to scalable, robust deployments will require careful focus on operational factors that far outweigh the initial capital expenditures of equipment and application software. Being vigilant in ensuring interoperability and being secure and flexible in provisioning and ongoing orchestration will ultimately lead to the optimal outcome and unleash the power of edge computing.

content provided by:



ZEDEDA provides a simple and scalable cloud-based orchestration solution that delivers visibility, control, and security for distributed edge computing, giving customers the freedom to deploy and manage any application on any hardware at hyperscale while connecting to any cloud or on-premises system.



Visibility, Control and Security for the Distributed Edge

ZEDEDA has architected its solution to meet the security, safety, uptime and usability needs of both OT and IT organizations, enabling them to focus on driving business outcomes. It is optimized to address the unique requirements for deploying computing at the distributed edge – outside of secure data centers, both on-premises and in the field. The solution leverages the open source EVE-OS from the Linux Foundation's LF Edge organization to provide an open, flexible and secure foundation while abstracting the complexity of the diverse hardware, connectivity and software at the distributed edge and eliminating any vendor lock-in.

Distributed edge solutions require a diverse mix of technologies and domain expertise, and ZEDEDA enables customers with an open, vendor-agnostic orchestration framework that breaks down silos and provides the needed agility and future-proofing as they evolve their connected operations. With ZEDEDA, customers can now seamlessly orchestrate intelligent applications at the distributed edge to gain access to critical insights, make real-time decisions and maximize operational efficiency.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.