



Building a Sustainable Enterprise Edge

Commissioned by



Executive summary

As enterprise data volumes grow exponentially and infrastructure becomes more distributed and disaggregated, a clear strategy at the edge is critical for IT and OT success. While hyper-distributed infrastructure has many benefits (lightweight, decentralized, low latency), there are also unique sustainability angles to consider. Although many organizations have taken an interest in improving the sustainability of their legacy infrastructure, there needs to be more emphasis on ensuring sustainability at the rapidly growing enterprise edge, a location that increasingly supports the most mission-critical, net-new digital workloads and applications.

To fully calculate the impact of edge compute on corporate sustainability goals and targets, enterprises must factor both “deploy the edge” and “use the edge” scenarios into their calculations.

For the first consideration — how best to build edge compute capacity with sustainability in mind — enterprises will encounter a few advantages, but an equal number of challenges. New edge locations can incorporate modern power and cooling approaches, including leveraging more renewable energy sources and more efficient compute. Processing data in place can also minimize data transport, helping network and cloud partners to improve their decarbonization story, thereby reducing an enterprise’s Scope 3 emissions. At the same time, more distributed edge infrastructure lacks the cloud’s critical mass and centralized sustainability benefits — a reality that must be factored into decision-making.

As for “use the edge,” edge- and IoT-enabled use cases can play a key role in helping enterprises reach their sustainability goals by improving efficiency, optimizing processes and delivering actionable, data-driven sustainability insights. About a quarter of enterprises today cite sustainability and/or ESG metric tracking as one of the drivers of their edge and IoT initiatives. While leveraging technology to support sustainability goals is in its relative infancy, IT- and edge-enabled sustainability approaches are likely to grow in importance as corporate strategies further align with evolving ESG mandates and targets.

Key findings

- **Enterprise data volumes and related power requirements are growing exponentially.** Data volumes for low-latency, edge-critical workloads are projected to grow at 80% CAGR through 2027, with power demand growing at 87% CAGR.
- **The edge will handle a good share of those compute requirements.** Over the same time frame, 62% of data volume and 68% of power demand will come from on-premises or nearby edge venues (with the remainder in core datacenters and/or cloud).
- **Enterprises must consider edge versus cloud venue decisions as part of their sustainability planning.** For example, newly built edge locations can leverage modern datacenter technologies to become more “green.” However, the more distributed nature of edge venues presents unique sustainability challenges and opportunities versus more centralized cloud infrastructure.
- **Edge-enabled use cases can contribute significantly to sustainability improvements.** More optimized, edge-enabled manufacturing lines can run more efficiently; smart buildings can analyze lighting, energy usage and other data locally to become more sustainable.
- **Corporate sustainability planning is in its early stages, as is the use of edge to contribute to those efforts — creating significant upside opportunities.** Just 43% of enterprises have a formal target in place for reducing environmental impact, with another 52% still in planning stages. The use of edge/IoT to directly impact sustainability is also in its infancy — just 25% of enterprises cite sustainability as a driver of their edge/IoT deployments.

The emergence of edge compute

Edge computing broadly refers to compute resources operating close to where data is generated, and outside of centralized locations such as hyperscale cloud. Today, the edge has evolved to include a variety of venues and form factors. On-premises can include anything from single-system servers or gateways to in-house micro-datacenter installations. But the edge today also includes near-premises venues offering edge as a service, such as via regional datacenter/colocation facilities or telecom-delivered multi-access edge computing (MEC) venues.

As enterprises collect more of this valuable edge data and leverage it to feed new mission-critical digital use cases, they face the important question of where best to run those workloads. Increasingly, much of that data is best stored, processed and analyzed where it originates — in a nearby edge venue. In the majority of edge use cases, relying on networks outside the premises presents an unacceptable availability risk. In some cases, sending that data to the cloud is too costly or results in performance or latency issues. In other cases, the data can't leave the edge due to data sovereignty or privacy concerns, a lack of available bandwidth or other restrictions.

Regardless of the driver, edge computing represents an increasingly important infrastructure layer to support the requirements of maturing digital transformation projects and the vast volumes of data generated daily in many enterprises. This brings the edge to a range of environments across industries and use cases. It could be a server closet in a hospital hallway storing patient data locally to minimize risk, an on-premises edge device on a manufacturing production line to reduce latency on mission-critical applications, or emerging MEC infrastructure to handle the real-time data demands of increasingly connected vehicles. In the wake of this growing demand, enterprise adoption of edge has become an increasingly mature, near-term reality:

- In 2023, 56% of enterprises have edge computing infrastructure “in use,” and another 28% are in trial or proof of concept, according to 451 Research’s Voice of the Enterprise: Edge Infrastructure and Services, Sourcing 2022 survey.
- Spending at the edge is slated to grow even further as 79% of enterprises are planning to increase spending on edge in 2023 — 32% of those “significantly” so, according to the same survey.


Increased edge adoption and spending is driven by massive increases in data growth — and recognition of the business and economic value that data can generate.

For instance, a 451 Research Market Monitor analysis of low-latency workloads across eight key industries in the US forecasts total data volume to grow at an 80% compound annual growth rate (CAGR) through 2027, from 5,700 PB in 2021 to 194,000 PB. In 2022, 52% of that volume was processed in a core or cloud location; that figure is slated to fall to 38% in 2027, as a range of on-premises and near-premises edge venues take on more processing load (see Figure 1).

Breaking that down a step further, while overall low-latency workload data volume is forecast to grow at a CAGR of 80%, growth at the edge is forecast to average just under 100% per year for that period.

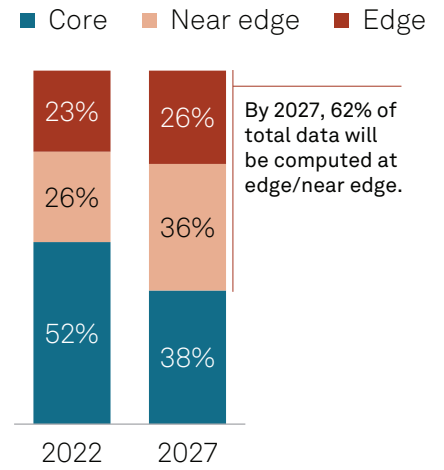
Figure 1: Total data compute volume and split by venue, low-latency workloads, US, 2021-2027

Total data compute (PB) initial & beyond initial analysis by venue segment



CAGR 2021-2027
 Core: 65%
 Near edge: 102%
 Edge: 91%
 Total: 80%

Total data compute (PB) % split by venue segment




By 2027, 62% of total data will be computed at edge/near edge.

Source: 451 Research's Edge Workload Total Addressable Market (TAM) Analysis, 2023.

That massive growth in data volumes creates another challenge: outsized power requirements to run the servers and datacenters needed to handle all of that data. According to the same 451 Research analysis, new digitally driven, data-heavy workloads will increase power demand by 87% CAGR in the same period, from 34 MW to more than 1,400 MW, with 68% of total power demand required in edge-type locations by 2027 (see Figure 2).

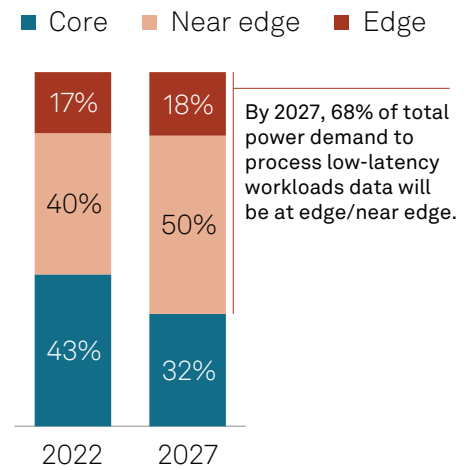
Figure 2: Total power demand and split by venue, low-latency workloads, US, 2021-2027

Total power demand (MW) for initial & beyond initial analysis by venue segment



CAGR 2021-2027
 Core: 71%
 Near edge: 104%
 Edge: 92%
 Total: 87%

Total power demand (MW) % split by venue segment



By 2027, 68% of total power demand to process low-latency workloads data will be at edge/near edge.

Source: 451 Research's Edge Workload Total Addressable Market (TAM) Analysis 2023.

Assessing the power and energy impact of all that data processing is a complex calculation, encompassing not just the primary compute infrastructure and associated power load, but also secondary considerations that impact enterprise sustainability calculations. For starters, large datacenters have immense cooling requirements, which put a strain on local water systems. Data moves from place to place as well, backhauled and transported in and out of telecom networks and hyperscaler cloud locations — data transmission that today is handled by telecom facilities that include power-consuming radio networks and datacenters in their own right. Enterprises consuming cloud services and sending data to and fro must also factor the sustainability impact of their partners' infrastructure into their own decarbonization targets.

Enterprise sustainability and the edge

Environmental, social and governance (ESG) concerns have risen to near the top of board-level agendas. On the world stage, the 2015 Paris Agreement on climate change and subsequent United Nations Climate Change conferences have brought government and, therefore, legislative attention to reducing carbon emissions by 50% across the globe by 2030. This is being pushed into the legal frameworks of nations. For example, the European Union's Green Deal, approved in 2020, not only sets out emissions targets, but also begins the process of establishing legislation on the circular economy, building renovation, biodiversity, farming and innovation. The circular economy and innovation aspects are fueling corporate change in IT.

Those mandates have made sustainability a corporate priority across an array of sectors. The World Economic Forum's Global Risks Report for 2023 cites failure to mitigate and failure to adapt to climate change as the top two risks over the next decade. However, only one in five companies has a plan in place to adapt to the physical risks of climate change, according to a [February 2023 analysis](#) of S&P Global Sustainable1 data.

For the purposes of this white paper, sustainability is measured by Scope 1, 2 and 3 emissions as defined by the Greenhouse Gas Protocol. Enterprise-owned datacenters or facilities with on-site power generation fall under all three scopes: Scope 1 for the energy directly generated, Scope 2 for the energy consumed but not generated and Scope 3 for the carbon emissions resulting from the construction of materials and systems, as well as their decommission, refurbishment, reuse, recycling or disposal. Cloud and multi-tenant datacenter providers, meanwhile, are also responsible for all three scopes of emission; however, their clients will typically categorize their share of this as Scope 3 emissions (corporate value chain).

As enterprises consider where to place their IT workloads — especially net-new digital workloads — they must calculate the performance and costs/benefits of edge versus cloud venues. Workloads typically better suited to remain at the on-premises edge or at a nearby edge service venue fall into three categories:

- Applications that require low latency and/or high bandwidth, such as extremely low-latency, high-frequency trading systems in finance or high-bandwidth computer vision workloads on a manufacturing line. Beyond those technical considerations, lack of viable, in-country cloud options or high data egress costs may also keep workloads on-premises and at the edge.
- Legacy systems such as mission-critical, typically on-premises production control systems such as SCADA systems that are intolerant to any form of outage. Such workloads often run in local air-gapped networks or have other security-related considerations, or run non-standard protocols that make local data processing a necessary consideration.
- Legal requirements from governments and industry regulators that require either data residency or data sovereignty, restricting data movement into hyperscale cloud providers.

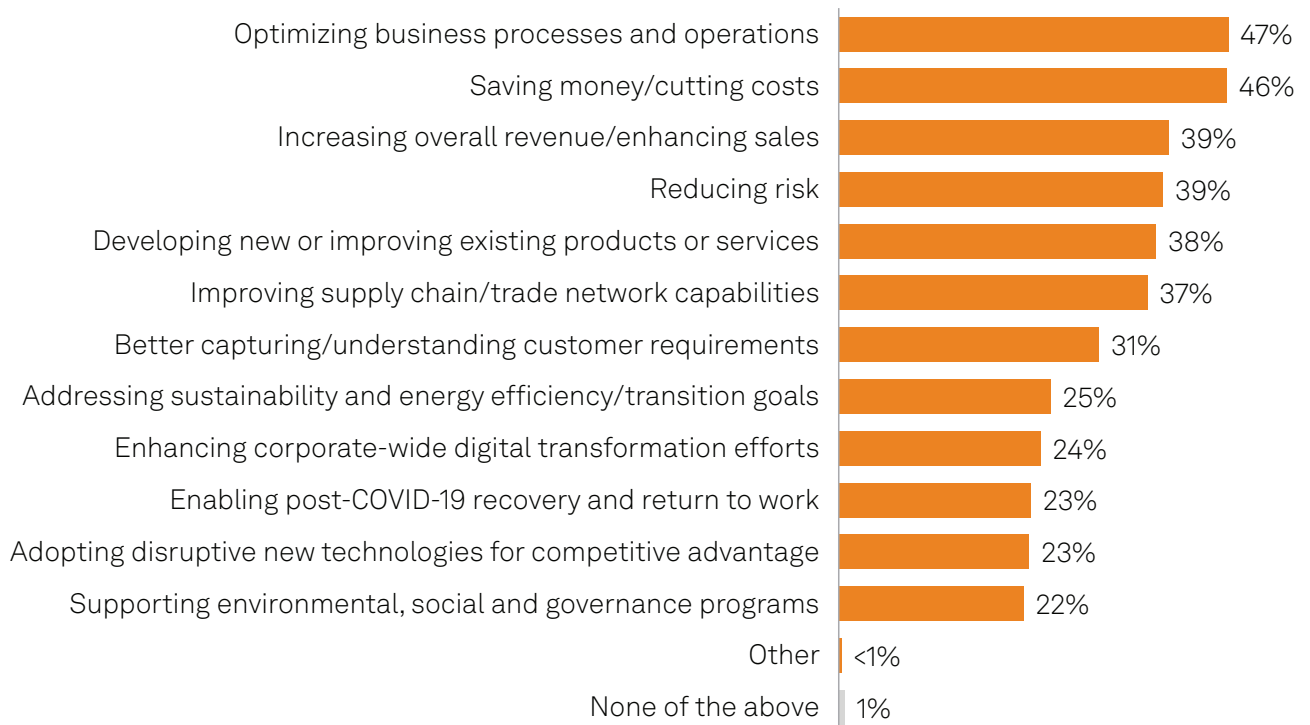
In many cases to date, such workloads have been housed in unconditioned, non-datacenter venues such as a retail back office or sitting just off the factory floor. When “the edge” is deployed in that way, it is hard to see a sustainability advantage when compared to the economies of scale and the power and cooling expertise of cloud and multi-tenant datacenter providers. That said, a more modern, well-thought-out approach to edge deployments — both on-premises and in nearby as-a-service edge venues — can deliver sustainability benefits in its own right. But it requires clear-headed planning:

- **On-premises edge locations could result in multiplication of infrastructure — adding to sustainability worries — if not planned properly.** For instance, mission-critical datacenter or nearby edge locations may require uninterruptible power supplies, on-site power generation, and redundancy in design, each of which brings sustainability challenges. Even if an edge venue doesn’t require this sort of infrastructure to start, it likely will as it supports more critical workloads. Indeed, refreshing existing technology and infrastructure in growing edge sites is the top driver of increased edge spending today, especially as such venues support more — and more business-critical — workloads, according to 451 Research’s Voice of the Enterprise: Edge Infrastructure & Services, Budgets and Outlook 2023 survey.
- **For enterprises with large and/or highly distributed edge processing needs, edge as a service offers a more sustainable alternative.** As local bandwidth improves and local datacenter/ colocation/MEC facilities proliferate, enterprises can turn to hosted edge compute options. Because such venues have many of the same advantages as larger cloud/datacenter locations (shared infrastructure, a degree of centralization, strong network interconnect capabilities, etc.), they are likely to deliver a more sustainable edge story than highly distributed, unconditioned on-premises edge compute. The evolution of edge compute infrastructure from purely on-premises to emerging as-a-service approaches will make it cheaper and more efficient to power than much more highly distributed on-premises deployments.
- **Telecom operators are building out edge compute venues as well, offering additional edge-as-a-service options for enterprises — and further potential Scope 3 sustainability gains.** Edge compute is an especially critical focus for telecommunications companies, especially when combined with emerging 5G network infrastructure. Indeed, in many instances, 5G represents a significant edge workload in and of itself, with virtualized network functions running not only in operator edge venues (such as at a distributed cell tower location) but also in hosted third-party edge locations such as with local datacenter partners. Edge/MEC services + 5G is of particular interest to enterprises requiring edge compute, bringing together the high speeds and ubiquitous coverage of 5G networks, the performance and latency benefits of nearby edge compute, and often fiber plant to backhaul data or send it off speedily to the cloud. Additionally, telecom edge/5G facilities may offer pass-through decarbonization benefits to enterprises because their shared facilities are likely denser, more centralized and thus likely more sustainable than highly distributed on-premises compute installations. In the end, telecom operators can benefit from leveraging their own edge deployments as well as helping enterprises both use the telecom edge and build/operate their own edges.
- **New edge infrastructure builds offer a significant sustainability opportunity: building new edge venues using modern datacenter technologies.** That includes new energy approaches such as using gas generators, free air cooling and offsetting heat through campus heating. Modernizing server infrastructure also brings potential efficiencies via newer, better-performing chips for the same energy consumption; more thought in overall server design to take advantage of new materials technology for cooling; and the use of relatively nascent cooling technologies such as liquid or immersion cooling.

Edge-enabled insights drive enterprise-wide sustainability benefits

Edge is one of the enablers of IoT, bringing data ingress, data processing and decision-making to a location that is not constrained by latency and bandwidth considerations or high cloud costs. IoT and edge-enabled sustainability efforts are early in their adoption curve, but they could leverage data insights to optimize production processes across a range of sectors — manufacturing, utilities, oil and gas, transportation, etc. — and impact overall building and site energy efficiency in the process. Projects directly impacting sustainability and ESG goals are still in their relative infancy — just 25% of enterprises say their IoT projects today are aimed at addressing sustainability or energy transition goals, while 22% cite specific IoT-enabled programs to support ESG efforts, according to 451 Research’s Voice of the Enterprise: Internet of Things, the OT Perspective, Use Cases and Outcomes 2023 (see Figure 3). However, as sustainability becomes a greater corporate priority, we expect those numbers to grow accordingly.

Figure 3: Sustainability, energy transition and ESG support are among key industrial IoT drivers



Q. Which of the following are drivers of IoT initiatives at your organization? Please select all that apply.

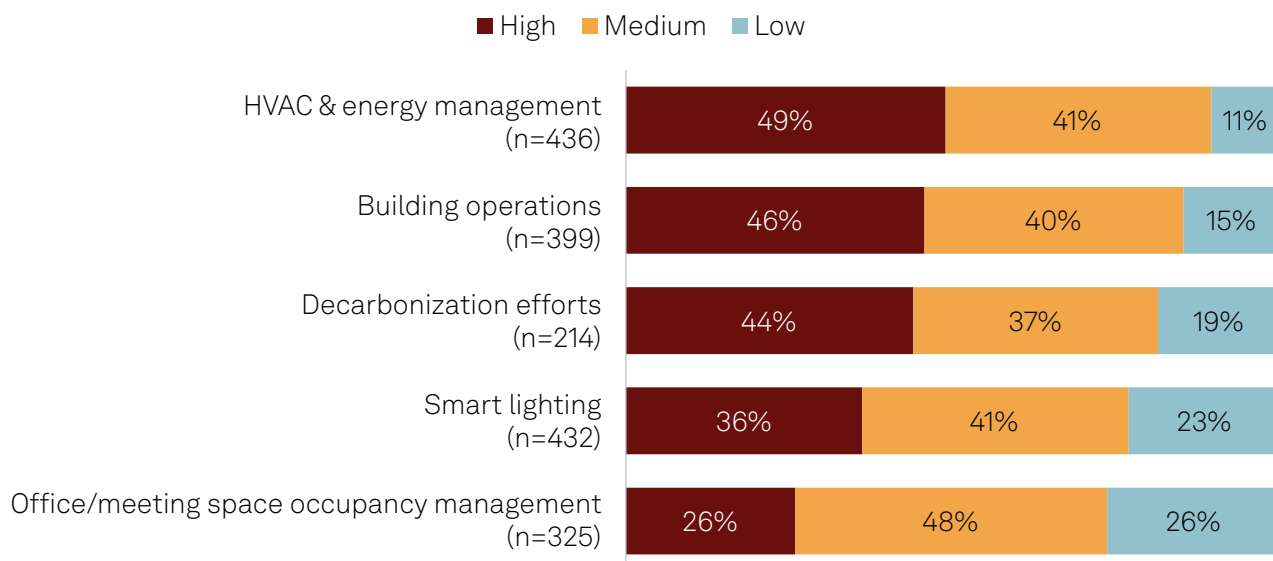
Base: All respondents (n=588).

Source: 451 Research’s Voice of the Enterprise: Internet of Things, the OT Perspective, Use Cases and Outcomes 2023.

Other use cases indirectly — but impactfully — affect sustainability, including leveraging edge and IoT technologies to optimize business processes and operations and reduce risk (compliance, safety, data protection, etc.). Creating more sustainable operations is a key industrial optimization goal, whether it be an individual plant or factory or an entire supply chain. IoT and edge can also contribute to an overall corporate approach to compliance risk management for government-mandated ESG targets, including reducing the risk of breach or non-compliance.

While it can be difficult in these early days to come up with a definitive figure for cost savings from sustainability at the edge, we can clearly see the impact that edge and IoT technologies can have on corporate sustainability efforts. For example, consider the effort to improve the sustainability of buildings — from office complex to factory venues to warehouse locations. Several sustainability use cases, in particular enabling building decarbonization and optimizing overall building operations, are expected to have a significant impact on enterprise building and site sustainability efforts (see Figure 4).

Figure 4: Enterprises expect IoT/edge-enabled use cases to have a significant impact on building sustainability



Q. Impact on overall business.

Base: All respondents (n=601).

Source: 451 Research's Voice of the Enterprise: Internet of Things, OT Perspective, Use Cases and Outcomes 2023.

Cost savings from controlling energy usage not only have a direct impact on facilities budgets but can also help fund other sustainability initiatives that take longer to deliver.



Implications

While edge computing is moving out of the nascent stage and into deployment in many industry verticals, we are only starting to see and understand edge's impact on sustainability. Building more efficient, "greener" edge infrastructure — be it self-managed on-premises or near-premises, or consumed via a growing number of edge-as-a-service providers — represents one important step forward. Beyond that, the edge's ability to deliver insights that can transform and optimize key business processes can also contribute significantly to strengthening an organization's sustainability stance.

1. [Siemens Smart Building Infrastructure, Desigo CC](#): Creating High Performance Buildings with Edge Solutions

For example, facility managers within the education sector can sync classroom scheduling software to their building's systems. By knowing when a classroom is due to be occupied, Desigo CC can turn on the lights, reset the temperature, disable any motion detectors, or adjust the blinds to the preference of the occupants ahead of their arrival. Once the class is over, it locks the door and returns the room to unoccupied mode, resulting in lower overheads from minimalized maintenance, improved security, and reduced energy usage. Desigo CC also helps clients to meet their sustainability goals by using demand response programs. By taking energy from local power systems during lower demand periods and returning it if the grid experiences peak demand, energy costs and carbon emissions are cut.

2. [ATOS case study](#): Saved 480 metric tons of CO2 emissions by automating processes at the edge

“For a retail customer, we’re talking about 66,000 hours saved in maintenance and compliance for maintaining the edge environment, which translates into about 480 metric tons of CO2 saved every year — thanks to automation and end-to-end monitoring.” Arnaud Langer, Global Edge & IoT Senior Product Director at Atos

3. [Nature Fresh Farms](#)

Being a non-GMO greenhouse, sustainability is a priority for Nature Fresh Farms. To aid its sustainability efforts, the company required an edge solution that could support AI and automation in temperature control, power management and water irrigation efficiency.

About the authors



Tiny Haynes

Principal Analyst for Edge Infrastructure and Services, Global

Tiny Haynes is Principal Analyst for Edge Infrastructure and Services, Global. He has been covering datacenters, cloud and edge computing, sustainability and data sovereignty for the past 12 years. Before that he was a product manager for 10 years, responsible for building and managing product portfolios for colocation and managed hosting.



Rich Karpinski

Principal Analyst, IoT and Industry Transformation

Rich Karpinski is Principal Analyst, IoT and Industry Transformation, leading the S&P Global Market Intelligence 451 Research IoT and Applied Infrastructure teams. In this capacity, Rich tracks, analyzes and anticipates the pace and direction of IoT and digital technology adoption across a range of industries and use cases. He also oversees 451 Research's quarterly survey of IoT adopters and twice-annual survey of operations technology (OT) professionals.

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About S&P Global Market Intelligence

At S&P Global Market Intelligence, we understand the importance of accurate, deep and insightful information. Our team of experts delivers unrivaled insights and leading data and technology solutions, partnering with customers to expand their perspective, operate with confidence, and make decisions with conviction.

S&P Global Market Intelligence is a division of S&P Global (NYSE: SPGI). S&P Global is the world's foremost provider of credit ratings, benchmarks, analytics and workflow solutions in the global capital, commodity and automotive markets. With every one of our offerings, we help many of the world's leading organizations navigate the economic landscape so they can plan for tomorrow, today. For more information, visit www.spglobal.com/marketintelligence.

CONTACTS

Americas: +1 800 447 2273

Japan: +81 3 6262 1887

Asia Pacific: +60 4 291 3600

Europe, Middle East, Africa: +44 (0) 134 432 8300

www.spglobal.com/marketintelligence

www.spglobal.com/en/enterprise/about/contact-us.html

Copyright © 2023 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 keeps certain activities of its divisions separate from each other 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 nonpublic 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.