

Building Hybrid Cloud Infrastructure for the Decade Ahead



Pathfinder

June 2021

Commissioned by



451 Research

S&P Global
Market Intelligence

About the Author



Eric Hanselman

Principal Research Analyst

Eric Hanselman is the Principal Research Analyst at 451 Research, a part of S&P Global Market Intelligence. He has an extensive, hands-on understanding of a broad range of IT subject areas, having direct experience in the areas of security, networks, application and infrastructure transformation and semiconductors. He coordinates industry analysis across the broad portfolio of 451 Research disciplines, contributes to the Information Security and Cloud Native Channels, and is a member of the Center of Excellence for Quantum Technologies.

The convergence of forces across the technology landscape is creating tectonic shifts in the industry, including 5G, SDN/NFV, edge computing and DevSecOps. Eric helps 451 Research's clients navigate these turbulent waters and determine their impacts and how they can best capitalize on them. For more than 20 years, Eric has worked with segment leaders in a spectrum of technologies, most recently as CTO of Leostream Corporation, a virtualization management provider. Prior to that, Eric guided security offerings for IBM and Internet Security Systems. At Wellfleet/Bay Networks and NEC, he was involved in the introduction of many new technologies ranging from high-performance image analysis to rollouts for IPv6.

Eric holds a patent in image compression systems. He is also a member of the Institute of Electrical and Electronics Engineers (IEEE), a Certified Information Systems Security Professional (CISSP) and a VMware Certified Professional (VCP), and he is a frequent speaker at leading industry conferences. Eric majored in Chemistry at Reed College.

Executive Summary

The dynamic nature of most business environments can place conflicting requirements on IT teams. They're working hard to find the resources to keep their businesses competitive, but they have to manage costs at the same time, all while supporting an increasingly diverse set of applications. More organizations are deploying hybrid cloud architectures to achieve these goals, but not all of them are succeeding. Many face challenges in bringing their on-premises environments up to the same level of agility they can achieve with their off-premises cloud resources. That imbalance can cause them to make trade-offs in workload placement, creating additional complexity in IT operations and limiting application development teams. Ensuring that the on-premises side of a hybrid environment is built on a foundation that can support the full range of possibilities that the modern business needs can be the differentiating factor in determining a successful outcome. Today's infrastructure has to be able to support a much greater scale, velocity and diversity of applications to stay competitive.

Key Findings

- Modern infrastructure has to support a complex mix of applications.
- Competitive pressures require organizations to manage capacity flexibly and cost-effectively.
- Hybrid operational efficiency requires coordination across resource realms.
- On-premises infrastructure has to be able to perform as well as cloud to make hybrid models successful.
- Many are making the hybrid journey without a formal strategy, increasing complexity and cost.
- Reducing operational complexity is critical to achieving management effectiveness and efficiency for hybrid models.

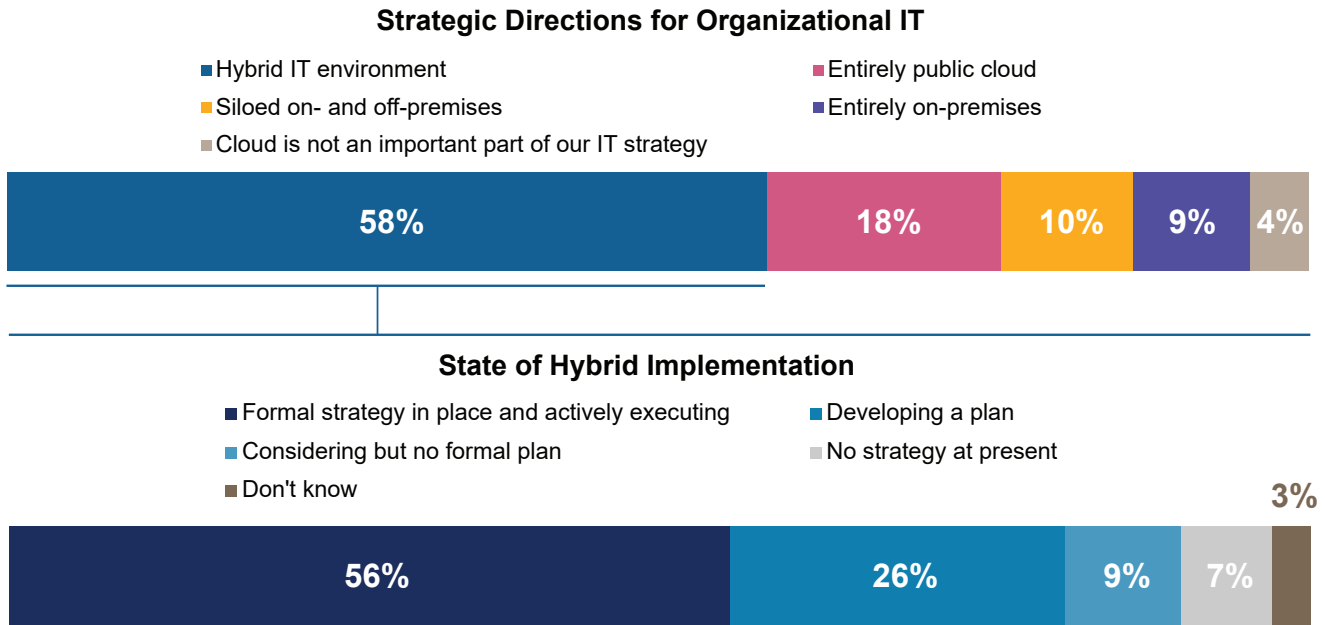
Introduction

The modern business is operating in markets that are buffeted by various forces that are difficult to plan for. Shifts in customer preferences and competitive position can seem constant. In order to succeed in such a turbulent world, organizations need to be nimble, and their technology resources can ensure that success. IT teams have to deliver infrastructure that can adapt quickly and scale, and optimize costs while doing so. These requirements have driven many organizations to deploy hybrid cloud environments, and while such environments offer significant advantages, they also bring complexities that can raise costs and hamper operations if not managed effectively. If their hybrid efforts are going to succeed, organizations need to ensure that their on-premises resources don't become a less-attractive or less-capable option. On- and off-premises resources have to operate in harmony to support application development efforts that deliver greater numbers of applications with more varied requirements and do it at much greater speed. To achieve this will require a better core infrastructure and operational improvements.

The Accidental Hybrid

The first use of public cloud for many organizations was experimental. Small projects or individual efforts were trials to see what it was and how it could be used. Today, most organizations have a deeper understanding of public cloud and are putting it to work in some form. The next stage in this journey is mastering the ability to coordinate operations across cloud and on-premises environments, a hybrid cloud operational pattern. According to a recent 451 Research study, 58% of respondents said their existing or planned IT operating environment is hybrid. That speaks to the needs that are driving hybrid adoption, which are many.

Figure 1: Enterprise Operating Environment, Current and in Plan



Q: Which of the following best describes your organization's existing or planned IT operating environment?

Base: All respondents (n=445)

Q: Which of the following best describes the state of your organization's strategy regarding hybrid IT?

Base: Hybrid as the organizing principle of IT strategy (n=147)

Note: In last year's survey, we asked this question of all respondents, with 34% reporting 'active execution' on hybrid strategy. The comparable figure in this year's survey is 41%.

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Vendor Evaluations - Quarterly Advisory Report

The modern IT environment is built from a large collection of applications, services, partners and capabilities. A web environment, for example, can include not only the content that's produced, but an ecosystem of technologies from many providers that bind together product listings, e-commerce functionality, marketing data, customer relationship management, inventory systems and security tools, among other things. Where that might once have all resided in an on-premises datacenter, it is now provided through on- and off-premises infrastructure. This shift has driven an expectation of hybrid operations and is changing IT operational models. The difficulty behind it is that these changes are often not fully planned out.

The Voice of the Enterprise (VoE) study also asked those respondents who said they were operating in a hybrid mode about the state of their hybrid implementations; the result is concerning. While there was improvement from the previous year's results, there are still 42% who said that they don't have a formal plan in place to guide their hybrid use, which indicates that they're arriving at hybrid more by accident than by design. This could be due to the incremental nature of the growth of cloud use – each new element is added for a specific project or requirement, but not coordinated through a larger plan. However, deploying hybrid cloud without a plan is a recipe for uncontrolled costs and escalating administrative work.

Hybrid Cloud Infrastructure Mindsets Have to Change

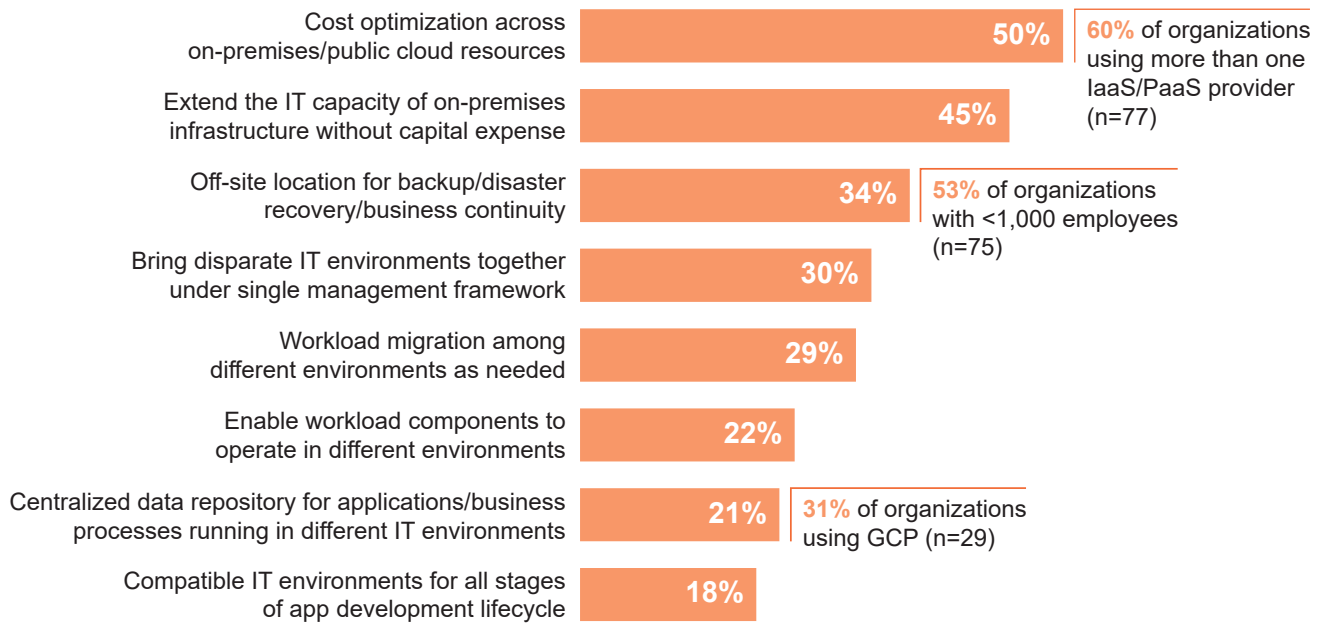
It's not hard to see where traditional operating mindsets contribute to the problems that result in uncontrolled hybrid operations. IT teams used to build infrastructure around the applications that it was going to support. Large, monolithic applications came with hardware compatibility lists that specified what was required. Virtualization started to change that situation years ago, but the mindset and its effects persist for some. Continuing to build infrastructure with a project-only focus, rather than planning for infrastructure that can support an organization's broader business goals, risks creating complexity and increasing costs. IT teams need to consider how the infrastructure that they're building will serve them over the long haul and how adaptable it will be to changing business and technology needs.

Adaptability is critical to infrastructure investment decisions. The mix of applications and development patterns is under constant change as organizations look to adapt and compete. Fundamental changes in building blocks, in the shift from virtual machines to containers and onward to functional and serverless models are already straining existing environments. As digitization moves forward, they'll need to support not only greater volumes of data coursing through their systems, but also provide the means to accelerate the processing and analysis of that data. Artificial intelligence and machine learning (AI/ML) need GPU and FPGA accelerators to scale to the heights required for rapid decision-making. All of this needs to be done with infrastructure that can flex and scale to meet these dynamic requirements.

Hyperscale cloud providers master adaptability with scale that enterprises are challenged to meet on-premises. They can deploy volumes of servers with a wide variety of technologies and capacities and refresh them at the same rate new technologies arrive. To reach the same level of adaptability without that scale, enterprises have to leverage systems that use modularity to let them reconfigure and integrate new technology with the same flexibility. It's critical for enterprises to be able to put new processing options, with next-generation processors, and increased networking capacity, like 400Gbit optical links, to work without causing major disruption to their datacenters. By disaggregating the traditional server model into resources that can be pooled and reformed as needed, they can support a wider set of use cases and application demands.

The shift to hybrid strategic thinking is being driven by all of these factors. That same 451 Research VotE study looked at use cases the respondents cited for pushing them forward to hybrid deployments. The leading use case was cost optimization across on-premises and public cloud resources (cited by 50% of all respondents and by 60% who indicated that they were using more than one IaaS or PaaS provider). That's a goal that is important to carry through to guide infrastructure decision-making. Creating different resource cost options with on and off-premises capacity is an important first step, but if they can't be managed efficiently, that cost advantage will be quickly consumed by the extra management overhead required to make it all work together.

Figure 2: Use Cases Driving Hybrid IT Deployment



Q: Which of the following use cases are driving your organization's implementation of hybrid IT environments? Please select up to three.

Base: Have or planning formal strategy for hybrid IT (n=125)

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Vendor Evaluations

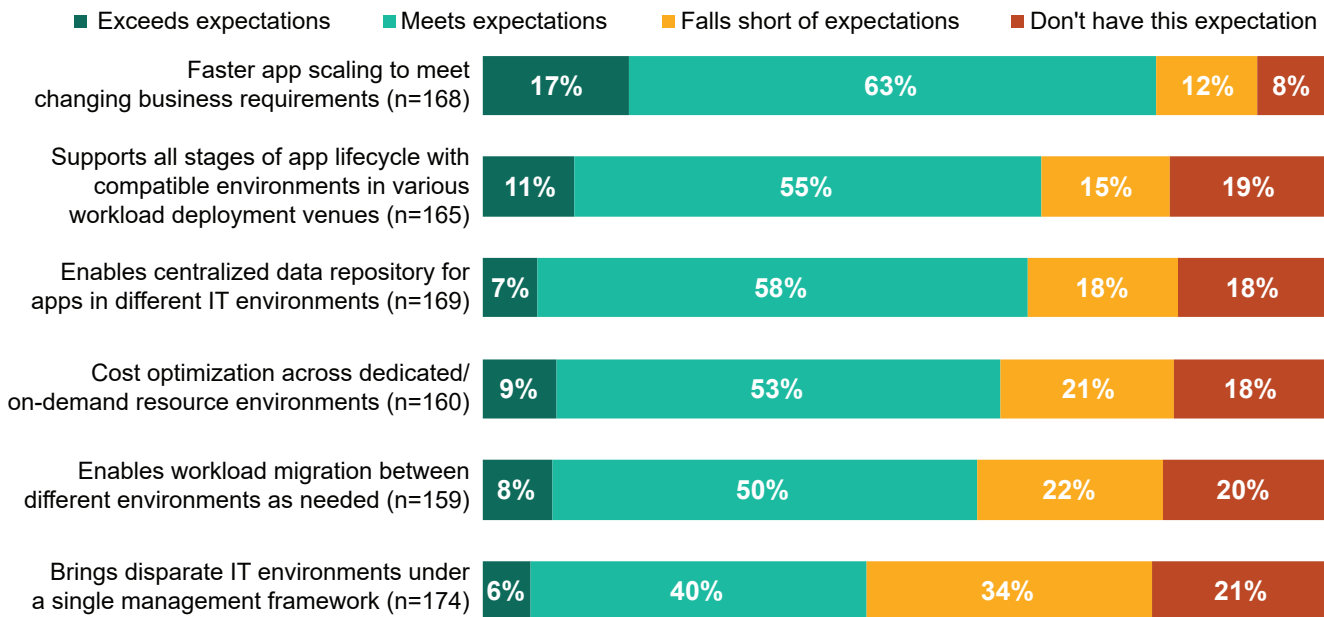
It's worth considering the state of enterprise infrastructure efficiency as part of the larger approach to hybrid. It's an area where expectations of scale and agility have been notably lower than those of cloud providers, but that is also often burdened with operational models that bear the legacy of traditional, segmented resource management. Without the ability to provision compute, memory, storage and networking in a composable fashion, an enterprise is limited in both flexibility and the speed and scale that it can accomplish.

Organizations' approaches to infrastructure management need to fully embrace cloud operational thinking in their on-premises operations. It's the next step in moving beyond the project-centric thinking that has held many back. It requires focusing on the flexibility of the provisioning process to deliver a wide range of resource options and be able to do it in a highly flexible fashion. To effectively match up with cloud resources, on-premises resources need to be able to operate at similar speeds and orchestrate with the same tools. Automation and orchestration have to be in place to deliver the velocity and scaling needed. They're technologies that can be force multipliers for IT administrators. They're also the technologies that can leverage newer, disaggregated infrastructure platforms that give enterprises the ability to assemble resources with the flexibility modern applications demand. Requiring application development teams to use different means to access different types of infrastructure locks in operational complexity that will slow them as well as the business imperatives they support and increase development time and costs.

Complexity Has To Be Tamed

The efficiency that hyperscale cloud providers have achieved comes from several aspects of their operations, but the most important is driving down management complexity. Enterprises need to keep this in mind as they look at what’s required to master hybrid operational models. The 451 Research VotE study asked respondents who have hybrid infrastructure in place how well it was delivering on their expectations. The vast majority were happy with the ability to provision quickly, but the lowest-scoring option was management simplification, likely because one of the casualties of heading into hybrid without sufficient planning – the accidental hybrid state – has been management integration. Even though cost optimization was the leading use case, it only scored in the middle of the pack for expectations. That’s an indication that enterprise goals aren’t always being met.

Figure 3: Hybrid IT: Delivery on Expectations



Q: For each of the following intended outcomes, please indicate whether your hybrid IT environment is meeting, exceeding or falling short of your expectations.

Base: Respondents with hybrid IT infrastructure in place

Source: 451 Research’s Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads & Key Projects

For organizations to achieve the cost and efficiency goals that hybrid offers, they need to reduce management complexity. It’s also a tactic to address chronic skills shortages in IT teams. By putting integrated management to work, they can reduce the need for specialized skills in each of the types of infrastructure they use, as well as put automation to work as a force multiplier for their administrative teams, and allow them to offer resources in the forms development teams want to consume them. Whether that’s containers delivered with Kubernetes or applications deployed with Terraform, IT teams can be ready to handle the demand.

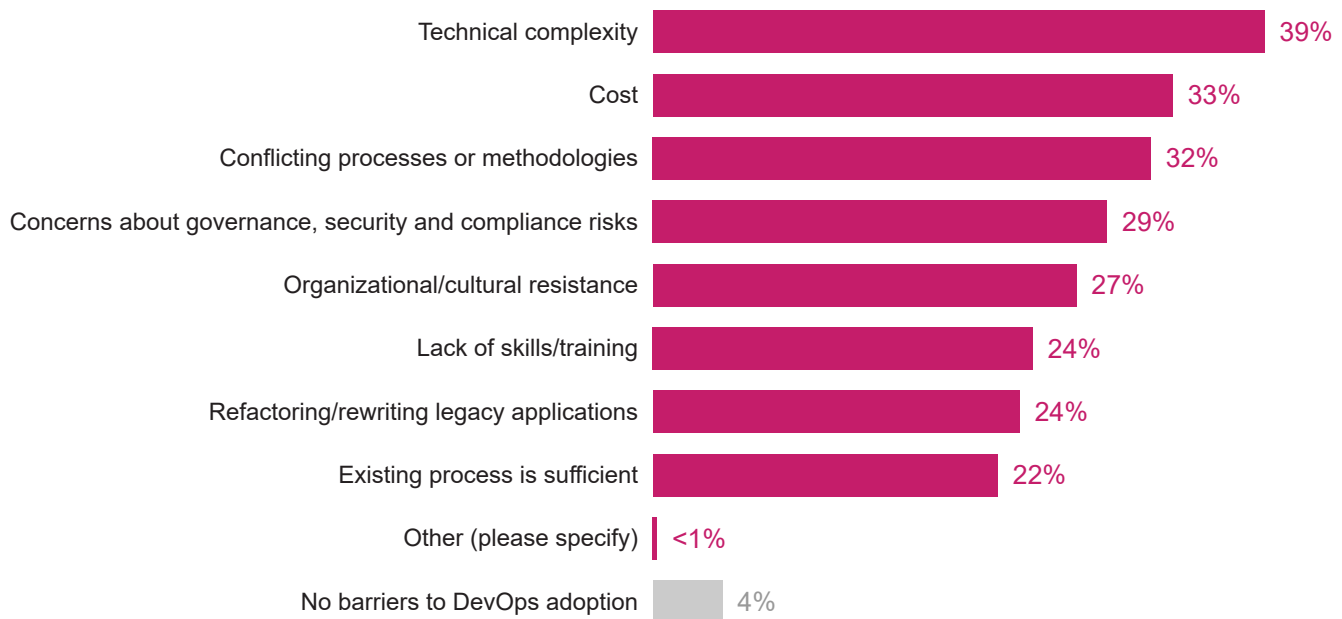
Management system deployment models place varying demands on administrative teams and on supporting infrastructure. The traditional models, where management systems were deployed on-premises, were geared toward the management of traditional infrastructure. They included a not inconsequential amount of operational irony. The systems that were meant to ease operational burdens came with their own operational burdens. They needed hardware to support them that needed to be managed, as well as requiring management themselves. As organizations look to management systems for hybrid environments, they need to consider the operational burden they'll place on the teams running them. To the greatest extent possible, management systems need to support an extension of the processes and procedures that are in place while integrating with new resources.

Management systems that are delivered in a SaaS model can reduce some of the operational toil required of administrative teams. There is limited work to deploy them, and the upkeep is managed by the provider. There is an additional benefit in hybrid infrastructure, in that they can have more performant and reliable access to cloud resources. SaaS models can also accelerate the velocity of innovation by enabling new functionality without the need for the disruption and downtime that traditional system upgrades require. They can be one more way to reduce operational complexity.

Putting Hybrid Infrastructure to Work

Getting to the level of infrastructure agility required to support the modern application development lifecycle isn't trivial, but it's mandatory for business success. Organizations struggle in a number of areas, but complexity is the predominant issue. New patterns of application development can be held back if existing infrastructure and operations are not up to the task. In a recent 451 Research VotE DevOps study, technical complexity was the leading reported barrier to advancing to a DevOps development pattern, one that's representative of more agile processes. Complexity was a greater concern than cost, indicating that spending won't fix the problem. Infrastructure has to be standardized, adaptable and automated to make it simpler for application development teams to consume. It has to be able to be reconfigured to fit new demands and new technologies, and it has to get to a cloud operating model to gain the speed and efficiency those teams demand. Hybrid infrastructure can offer the capabilities to meet those demands by delivering deployment options that suit the diverse needs of today's applications.

Figure 4: Top Barriers to DevOps Deployments



Source: 451 Research's Voice of the Enterprise: DevOps, Organizational Dynamics

Q: What are/were the primary barriers to deploying DevOps for your organization? Please choose no more than three.

Base: All respondents, abbreviated fielding (n= 357)

Organizations can achieve a fully functional hybrid infrastructure quicker by building a foundation that more easily supports the hybrid operating model, as well as the automation and orchestration capabilities to deliver speed and standardization that make IT teams more effective. When that foundation is meshed with management systems that support and extend to resources wherever they're needed, organizations can be ready to support the twists and turns that their markets and customers require.

Conclusions

Deploying hybrid environments can make organizations more competitive by accelerating their ability to deliver a more diverse mix of applications faster and at greater scale. But they require modern infrastructure and operational efficiency to get there. Organizations need to take the time to thoughtfully plan and invest in building systems and operational processes that will expand their opportunities, not weigh them down with complexity. Crafting on-premises infrastructure that can dynamically meet the ever-changing needs of modern apps and enable easy integration with off-premises resources is necessary to achieve this goal. By simplifying and standardizing their operations, organizations can expand their ability to innovate and unleash their potential. If they want to avoid becoming an accidental hybrid, with uncontrolled costs and limiting operational constraints, they have to break out of old patterns of thinking and move forward. They need to build environments that are engineered to operate for the long term.



Learn more about the future of hybrid cloud infrastructure at cisco.com/go/ucsx.

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 Web sites, 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.