

Mitigating Risks in the Hybrid Multicloud Journey



Pathfinder

January 2022

Commissioned by

kyndryl[™]

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 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 move to a hybrid multicloud environment is already a reality for some and may seem inevitable for many. That shift unintentionally brings a set of complexities that can strain traditional approaches to cyber resilience, availability and compliance. The natural expansion that is pulling enterprises into environments outside of their traditional datacenters is also taking critical application components and data beyond the protections that they've had in place. It's necessary to establish those protections in these new venues, but it can be resource-intensive and challenging for organizations that haven't had the time to develop the technical depth to do it effectively.

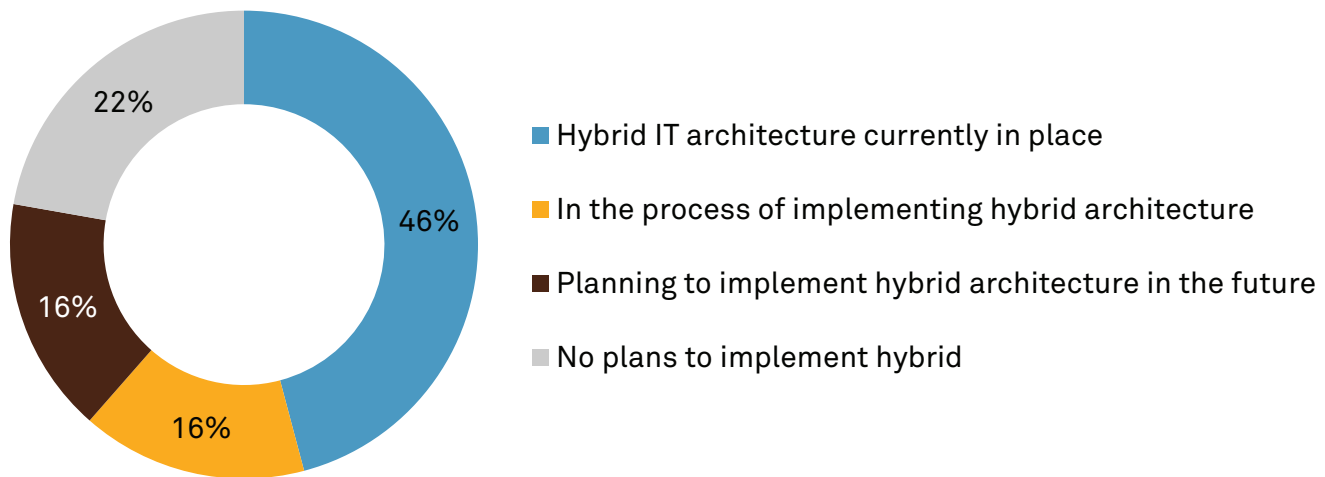
Key Findings

- Hybrid multicloud environments require new strategies to manage risks.
- Cyber resilience requires action now to mitigate risks to business.
- Changes in attack patterns and tools require protections for data that can mitigate new threats.
- Data protection in hybrid multicloud requires enhanced vigilance to new risks.
- Hybrid data management improvements are imperative in light of regulatory forces like the EU's General Data Protection Regulation and the California Consumer Privacy Act.
- Automation and orchestration are required to deal with the scale of hybrid environments effectively.

Benefits of Hybrid Multicloud

As digital transformation and cloudification continue to expand, enterprises are stretching their infrastructure to massive scale. Creating a hybrid multicloud environment should spike interoperability and allow for a breadth of partnerships and collaborations across clouds. According to recent 451 Research data, 46% of organizations already have a hybrid IT architecture in place, and a further 32% are actively planning or currently in the process of implementing one (see Figure 1).

Figure 1: The Future Is Hybrid for Most



Q: Which of the following best describes the current state of your organization's IT environment?

Base: All respondents (n=423)

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2021

Public clouds are gaining momentum as key workload deployment venues. The pivot to public cloud continues, with 32% of organizations pointing to IaaS/PaaS and SaaS as the primary venue for deployment of workloads and applications in 2021, up from 26% in 2020. We predict that hybrid environments will become too widespread for businesses not to have some level of hybrid collaboration, even if it's only for a few applications.

The growing use of cloud infrastructure creates an environment where countless configurations of virtual resources could be in use. A vast range of available resources means that in order to be successful, enterprises need to be able to create assets that can exist in multiple environments. A hybrid IT strategy allows more room for collaboration and interoperability.

In order to successfully operate in a hybrid multicloud mode, organizations have to be able to move data across infrastructure without sacrificing fluidity. As organizations increase their use of containers and microservice architectures, the lightweight application components themselves can be easily transferred across clouds. In an environment where data may need to be stored at the edge – such as an IoT deployment or a smart factory – it's imperative that that data can be easily transferred across clouds to allow for more efficient use of the distributed infrastructure.

The technologies to enable hybrid multicloud operation are available and in widespread use. Strategies for data movement and availability can be tailored to the needs of the applications they're supporting. Replication and publish/subscribe techniques offer basic approaches to ensure data is available where it's needed. More complex paths – distributed databases like MongoDB or Cassandra, for example – can span locations while automating the task of data distribution.

One more driver of the shift to hybrid is the expansion of an organization's technology ecosystem. There are two common paths: the embrace of an important technology and the presence of a partner or provider. In either case, the organization establishes a presence in a new venue to take advantage of the technology or service. Public cloud providers offer specialized technology, such as image or speech recognition and machine learning capabilities. To use the technology, data has to be available in that cloud environment, and the results are delivered there as well. Services offered by ecosystem partners, such as customer marketing or engagement, may be hosted with a particular provider, making it attractive to have application components hosted there for improved performance. All of these are factors that may push organizations into multiple environments where they need to manage the reliability and availability of data.

Complexities of Hybrid Multicloud

As cloud infrastructure expands, it becomes increasingly convoluted – opening the door for attacks, errors and failures. An enterprise operating in a hybrid multicloud mode may develop risk exposures in its infrastructure without being operationally conscious of their existence. This is a challenge that can develop with organic infrastructure growth. If the organization doesn't have processes to onboard new resources that reevaluate risk at each step, threats to resilience can creep in. It's not uncommon for the use of various cloud or hosting environments to be uncoordinated, and too many connected workloads can create attack vectors that are foreign to existing infosec teams and software.

Working across cloud environments also creates a particular dependence on interconnection, which has less than perfect reliability. Hybrid environments extend paths to data across typically dissimilar technologies and with inconsistent tools for managing and monitoring them. Managing key data paths can be challenging enough within a datacenter. Once that data is dispersed across an infrastructure at massive scale, it becomes an even steeper challenge.

One of the larger difficulties that interconnection presents is that the failure modes that are introduced can be much more complex. That can make detecting failures and recovering from them difficult. For example, a path that shares traffic from a number of sources can become congested, creating increases in latency or packet loss. For applications that depend on timely synchronization, increases in latency over their performance threshold can have an effect that's similar to the failure of the path, but yet to monitoring tools, they can still appear to be functioning. Diagnosing problems like this is problematic, particularly because they're often experienced only under significant load, which can make their occurrence intermittent.

An application's failure modes can be complicated by factors within the different environments in which the components are located. Local performance problems can be driven by a host of issues that range from application errors, storage I/O variability, instance sizing errors and simple failures. The complexity arises in trying to determine that a failure has occurred and then recovering – across environments where the detection and recovery mechanisms are unique. Organizations tackling this problem can wind up expending considerable resources to develop skills in each new environment in which they operate.

Hybrid IT infrastructure can also be a headache for operations teams that have to monitor more environments than ever before. While IT operations experts may be well versed in managing their company's private on-premises servers, in a hybrid multicloud configuration, they will need to be interoperable with public clouds, and potentially a private cloud from another provider. It's difficult to maintain operational efficiency when teams are tasked with mastering different skill sets and integrating the results into a working process.

One of the larger issues with hybrid environments is that underlying risks may be masked by the complexity of responsibility matrices and the application structures that are built across them. The combination of all of these factors can accrue a hidden set of potential problems that aren't taken into account in business continuity and disaster recovery planning, which looks at each environment independently.

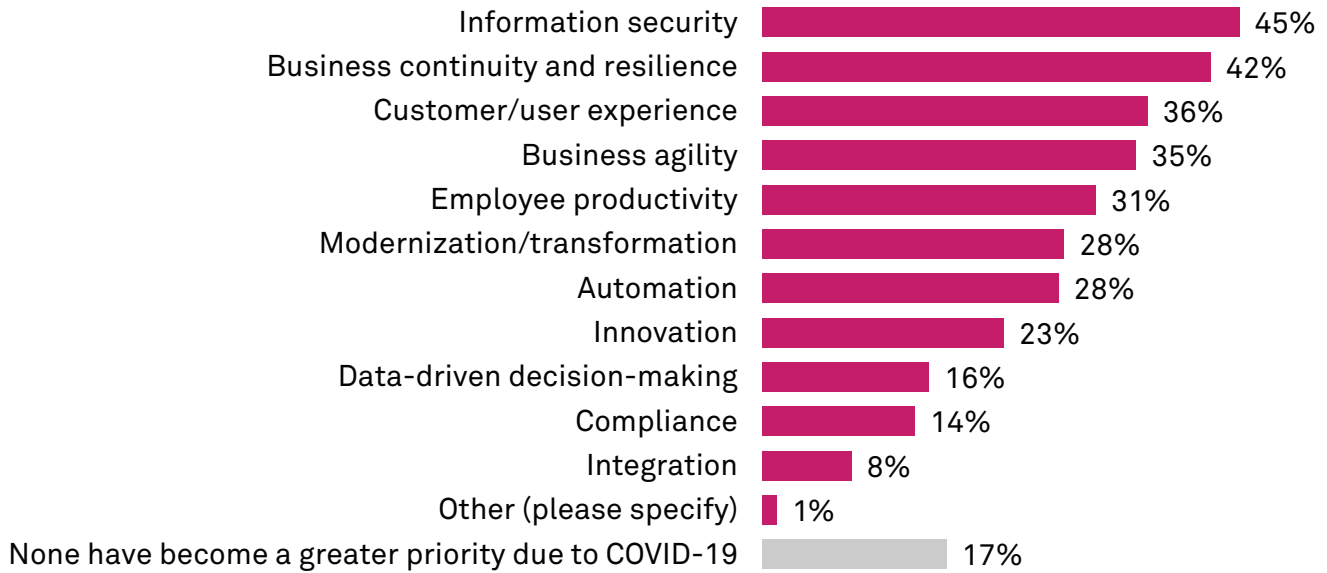
Resilience Imperatives

With the combination of ecosystem expansions and a set of benefits driving the adoption of hybrid environments, organizations have a strong imperative to address the resilience of this new environment to ensure that they can maintain the same levels of availability in key applications that they've had traditionally. This expansion trend is not a one-time event; it's a new reality. New environments will continue to offer value in new ways. Traditional IaaS clouds have given way to container environments, and serverless and functional environments are playing a more prominent role. This means that organizations have to create capabilities that make it simple to extend the protections needed for providing resilience to new services or execution venues.

This imperative has to be acted on today. It's not a matter of simply delaying a single project to make a single new environment resilient. Any delay defers the development of an important skill – the ability to leverage new infrastructure with confidence. Once in place, it unlocks the ability to support a nimble infrastructure strategy by ensuring that no matter how infrastructure needs are met, the robustness of the business services and applications running on it is ensured. There's a lot of discussion about the orchestration and automation that is needed for agile infrastructure, but agile resilience is just as important. A recent 451 Research study (see Figure 2) shows how the strains created by the global pandemic reinforced the importance of resilience: organizations gained a greater awareness as they grappled with change. Respondents cited business continuity and resilience as the second-highest technology objective, which has created new urgency for resilience imperatives.

There's a lot of discussion about the orchestration and automation that is needed for agile infrastructure, but agile resilience is just as important.

Figure 2: Prevailing Priorities for Enterprises Post-Pandemic



Q. Which of the following technology objectives, if any, have become a greater priority for your organization due to the influence of the coronavirus (COVID-19) outbreak? Please select all that apply.

Base: All respondents (n=371)

Source: 451 Research's Voice of the Enterprise: Digital Pulse, Coronavirus Flash Survey October 2020

There are a number of components to this imperative that organizations need to cover in order to effectively deliver resilience in ways that will be operationally efficient in hybrid multicloud deployments. Enterprises can address some of these components by expanding existing business continuity and disaster recovery to include partner resources. Most business continuity and disaster recovery planning exercises consider owned assets, which limits the extent to which organizations account for the capacity delivered by hosting providers or public clouds. Some of this is because, historically, this was complicated to achieve – most traditional business continuity and disaster recovery practices couldn't easily extend off-premises, and those that could required significant manual intervention. With appropriate automation and orchestration, on- and off-premises infrastructure can now have the same levels of protection.

Information security needs drive another significant component of the resilience imperative. Hybrid multicloud environments have a much larger attack surface. Because attackers are increasingly using automated tools, it has become much easier for them to find and target different elements of a more distributed application infrastructure. Another benefit of resilience capabilities that support multiple environments is that they make it possible to recover into infrastructure that isn't under attack. This can reduce the risk that any single element of the full business process implementation can take down an application.

Requirements for Resilience

Any resilience approach has to meet a set of requirements to provide the level of functionality necessary to support hybrid multicloud environments. First and foremost, the approach has to extend awareness and visibility across the entire hybrid environment. Having a common reference point that can act as a shared resource can bring teams together and provide a more complete perspective of the current state of an organization's infrastructure. To accomplish this, it has to span physical and virtual resources and provide equivalent perspectives. Across these realms, the resilience approach has to create service abstractions that can simplify operations by translating high-level capabilities into the native functionality in each environment. Approaches that require specialized knowledge for different domains can't scale and will make the process of onboarding new environments costly. Having common services in different venues has multiple benefits: application teams can deliver applications and services more quickly because they have to do little new adaptation, and it reduces the potential to be locked into a particular environment because of a reduction in dependence on environment-specific functionality.

Any approach also has to be flexible enough to work well within different operational environments. Having the agility to be deployed quickly can mean that establishing protections doesn't hold back experimentation or fast reactions to market changes. Scalability should be a natural byproduct of this level of agility. One of the main challenges of hybrid environments is scale.

One of the aspects of resilience capabilities that should be driving the support of larger scale is automation/orchestration. It's worth considering as its own requirement because it is such an important element of any implementation. Effective automation and orchestration should be the vehicle that delivers abstractions while reducing the workload of the operations team.


The timeliness of recovery and the breadth of recovery options are also important requirements. In many cases, the two go hand in hand because having more options for recovery may allow optimization of the recovery process to suit the needs of different situations. In hybrid environments, outages can have many factors that are intertwined, creating dependencies that can preclude certain recovery paths. Effective approaches will be able to offer alternatives to work around any blocking issues.

A resilience approach that addresses these requirements can make organizations more agile by allowing them to adapt more quickly and recover from problems faster.

Approaches to Greater Resilience

Hybrid multicloud environments have enough different facets that it can be difficult to identify how to ensure overall resilience. Deciding whether to extend existing data protections, capitalize on native services in new environments or take on wholly new methods isn't simple. To make detailed decisions can also require in-depth knowledge of the technical details of the various environments that may be outside of the skill set of IT teams. It's highly likely that they won't be skilled in cloud or hosting services that are new to the organization, and taking the time to develop those skills could either hold back new services and applications or leave open the possibility that operational risks won't be identified and mitigated. This is an area where a capable service provider partner can be particularly useful in both identifying issues and providing perspectives on how to deal with them.

Figure 3: Current Skills Shortages Needed to Manage Cloud Environments



Q: Which of the following skills categories are most acutely lacking when it comes to managing your organization's cloud environment? Please select all that apply.

Base: Current or future cloud users, abbreviated fielding (n=330)

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Organizational Dynamics 2021

Most organizations are struggling to keep up with skills to manage new infrastructure models. Respondents to a recent 451 Research survey said that cloud platform expertise is their most acute skills shortage, outpacing security by nine percentage points (41% to 32%), which is a perennial primary concern (see Figure 3). In situations like this, it can be difficult to hire and retain the necessary talent to meet operational needs, let alone staff teams to make strategic decisions. Working with a specialist service provider partner can extend the skills of existing staff and bolster them with service capabilities that can address the complexities that hybrid models present. A working partnership will allow organizations to achieve the necessary scale on their own terms.

This is a process that can have significant benefits for the organization. Putting comprehensive hybrid resilience capabilities in place can help organizations get ahead of the needs of their development teams. They can manage data resilience needs today, but, more importantly, they can provide a foundation that development teams can leverage over time, making developers less dependent on native and proprietary options that exist in individual cloud providers. Putting comprehensive hybrid resilience capabilities in place can also expand an organization's options for infrastructure choice, making it easier to optimize environments to fit the needs of the business. At the same time, it can enable organizations to respond more rapidly to changing market conditions and vendor relationships.

Conclusions

The shift to hybrid multicloud infrastructure models is well underway for many organizations. Such models offer benefits that can be compelling, and many organizations will drift to that mode of operation without fully considering its impact on the reliability and resilience of their application environments.

All organizations, particularly those that have yet to fully adopt this model, need to consider how they'll address the associated risks and manage them in an operationally efficient manner. Addressing this now can help them manage the current environment, as well as provide a means to confidently handle infrastructure expansion. It's a process whose value can be maximized by working with a capable partner that can deliver guidance in an area where there are often considerable skills gaps. Increasing application resilience in a hybrid world has many complexities, and it's a goal that is extremely valuable to achieve.



content provided by: **kyndryl**[™]

In a complex heterogeneous cloud environment, achieving operational resilience demands an integrated approach to detect and respond to threats, safeguard data, ensure high availability, and quickly recover critical business processes and systems in the case of a disaster. An integrated cyber resilience strategy and plan comprising the skills, services, and technologies that support growing compliance requirements is the first step in creating such an orchestrated platform.

Kyndryl Security & Resiliency provides a comprehensive range of platform-agnostic services to assist businesses in developing and implementing enterprise-wide cyber resilience policies to help them de-risk their journey to cloud. Kyndryl's expertise, processes, and technologies help enterprises keep their critical systems secure, available, reliable, and recoverable across heterogeneous environment regardless of their size and complexities, while supporting evolving compliance needs.

Kyndryl's Security & Resiliency services help clients across the world design, build, transform, and manage cyber security and resiliency capabilities – ranging from integrated threat management platform with detection, protection, response, and recovery to hybrid platform recovery with data protection, disaster recovery, high availability, cyber resilience service, and IT resilience orchestration services. For more information, please visit <https://www.kyndryl.com/services/business-continuity>

About Kyndryl

Kyndryl is the world's largest IT infrastructure services provider. The company designs, builds, manages, and modernizes the complex, mission-critical information systems that the world depends on every day. Kyndryl's nearly 90,000 employees serve over 4,000 customers in more than 60 countries around the world, including 75 percent of the Fortune 100. For more information, visit www.kyndryl.com.

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 © 2022 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.