451 Research®

PATHFINDER REPORT

Mitigating Risks in the Hybrid Multicloud Journey

RESILIENCE IMPERATIVES

COMMISSIONED BY

IBM

NOVEMBER 2019

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

CHIEF ANALYST

Eric Hanselman is the Chief Analyst at 451 Research. He has an extensive, hands-on understanding of a broad range of IT subject areas, having direct experience in the areas of networks, virtualization, security and semiconductors. He coordinates industry analysis across the broad portfolio of 451 Research disciplines. The convergence of forces across the technology landscape is creating tectonic shifts in the industry, including SDN/NFV, hyperconvergence and the Internet of Things (IoT). Eric helps 451 Research's clients navigate these turbulent waters and determine their impacts and how they can best capitalize on them. Eric is also a member of 451 Research's Center of Excellence for Quantum Technologies.

Executive Summary

The move to a hybrid multicloud environment may already be a reality for some and can seem inevitable for many. That shift unintentionally brings a set of complexities that can strain traditional approaches to availability, security 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 may not have had the time to develop the technical depth to do it effectively.

Key Findings

- Hybrid multicloud environments require new strategies to manage risks.
- 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 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2019 survey, 57% of organizations would describe their IT approach and strategy as a hybrid one.

Figure 1: The future is hybrid for most

Source: 451 Research, Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads and Key Projects 2019
Q: Which of the following best describes your organization's overall IT approach and strategy?

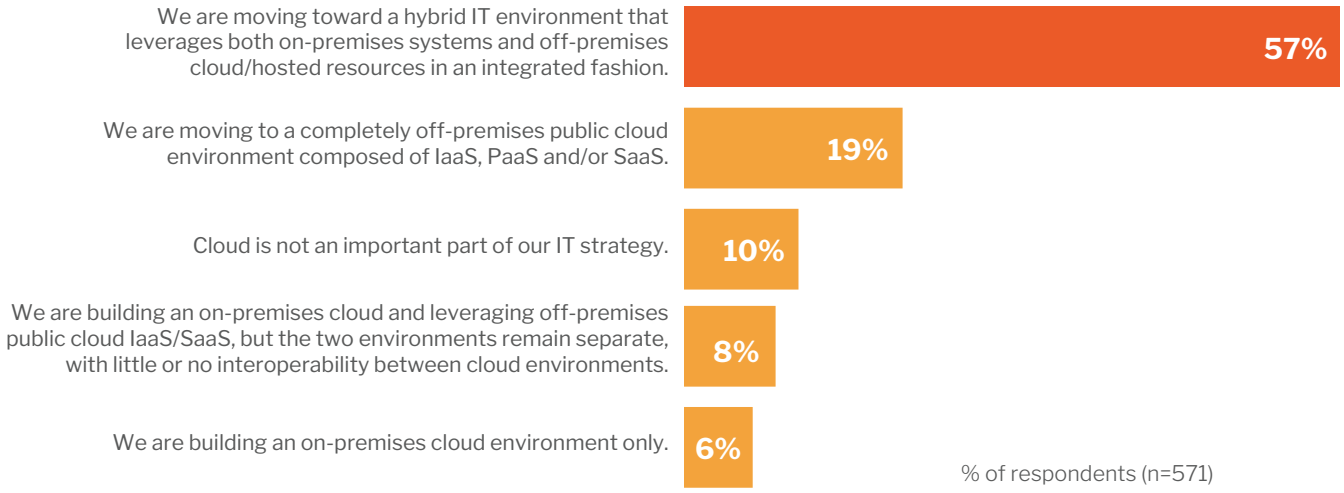

It may be difficult for the 19% of respondents that report moving toward a public cloud environment to exist exclusively in public clouds. We predict that hybrid environments will become too widespread for businesses to not have some level of hybrid collaboration, even if it's only for a few applications.

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 they move toward the use of containers and microservices 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, such as distributed databases like MongoDB or Cassandra, 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 them, 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 in 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 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.

Failure modes of an application 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. If an organization wants to tackle this problem, it can wind up having to expend considerable resources to develop skills in each new environment in which it operates.

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 risks of hybrid environments is that underlying risks may be masked by the complexity of 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 is deferring the development of an important skill that can support a nimble infrastructure strategy by ensuring that no matter how infrastructure needs are met, the robustness of the services and applications running on it are 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.

There are a number of components to this imperative that have to be covered in order to effectively deliver resilience in ways that will be operationally efficient in hybrid multicloud deployments. Some of them can be addressed 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 capacity delivered by hosting providers or public clouds is accounted for. 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.

Another significant component of the resilience imperative is driven by information security needs. Hybrid multicloud environments have a much larger attack surface. Because of the rapid increase in the use of automated attack tools by the attacker community, it has become much easier to find and target different elements of a more distributed application infrastructure. An additional 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

To provide the levels of functionality necessary to support hybrid multicloud environments, there is a set of requirements that any resilience approach has to meet. First and foremost, it 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, it 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 that application teams can expect in different venues has multiple benefits: Applications and services can be delivered more quickly because little new adaptation has to be done, and they reduce 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 operations team's workload.

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 would 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 2: Current skills shortages by IT category – 2019 and 2018

Source: 451 Research's Voice of the Enterprise: Digital Pulse, Organizational Dynamics Quarterly Advisory Report

Q: In which of the following IT categories, if any, is your organization currently facing an acute skills shortage? Please select all that apply.

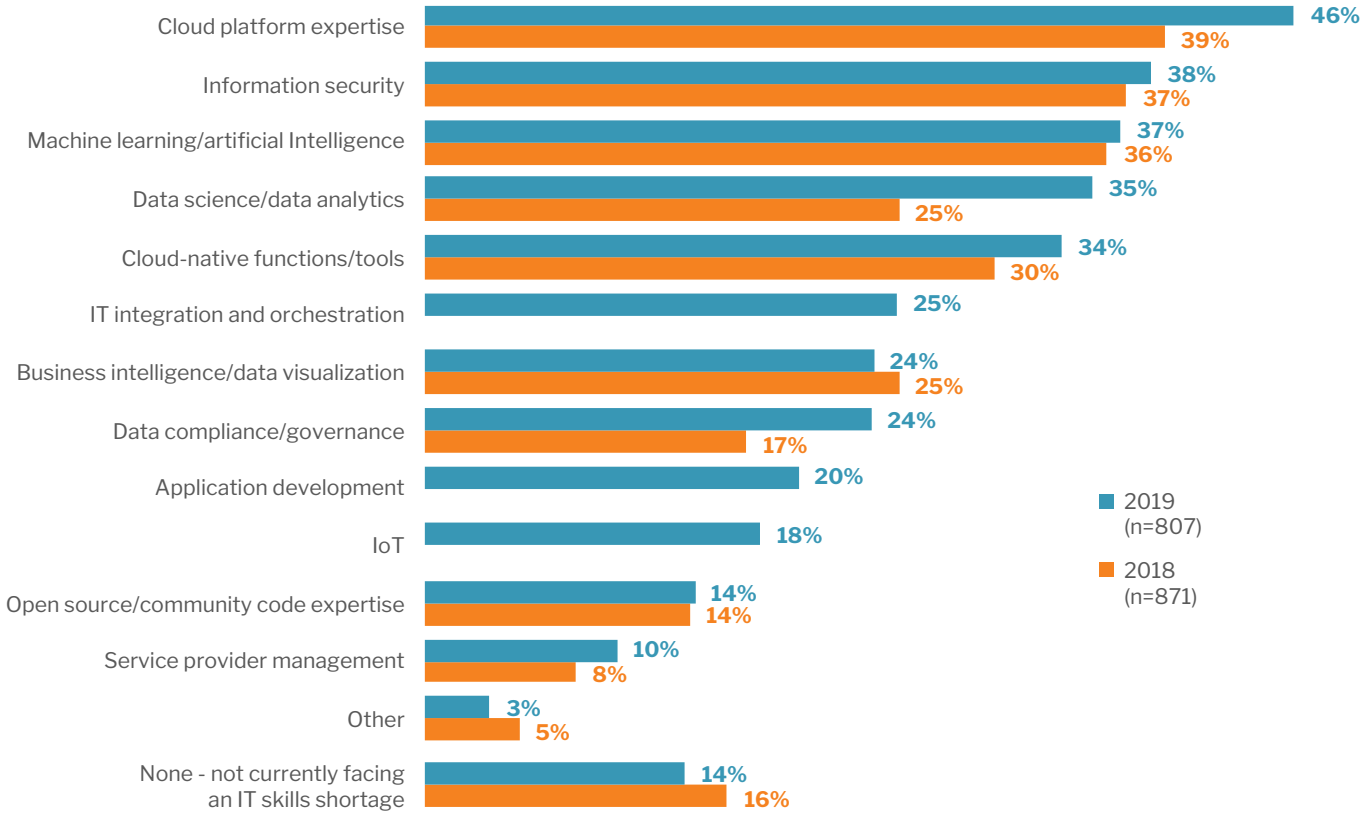

Most organizations are struggling to keep up with skills to manage new infrastructure models. In 451 Research's Q1 2019 Voice of the Enterprise: Digital Pulse study, respondents said that cloud platform expertise was their most acute skills shortage, outpacing information security (by 46% to 38%), which was the leader in earlier polls. 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. It can 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 and Recommendations

The shift to hybrid multicloud infrastructure models is well underway for many organizations. It offers 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 to 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.

Sponsor Profile

Achieving resilience in a complex hybrid multicloud environment requires an integrated platform to protect data, maintain high availability, and recover vital technology infrastructure and systems rapidly in the event of a disaster. Building such a platform starts with an integrated resilience strategy and plan encompassing technologies, business processes, people, and policies.

IBM Services helps clients develop and implement enterprise-wide resilience strategies and solutions to help de-risk their journey to hybrid multicloud. It helps clients optimize business and IT availability and continuity, either within day-to-day business and IT operations or under unexpected conditions such as cyberattacks, hardware and software failures, supplier failures, and natural or human-made disasters. It supports businesses across hybrid multicloud environments, including public cloud, private cloud, colocation and on-premises data center environments. It also has a strong multicloud practice across popular cloud providers, including Red Hat OpenShift, AWS, Azure, Google Cloud, and IBM Cloud.

IBM's portfolio of resilience offerings includes advisory, infrastructure, design/build, implementation and managed services – ranging from data protection, virtualization, disaster recovery, and cyber resilience to full-scale compute, data and application resilience, high availability, and efficient facilities and data centers. Using software-defined approaches, cloud-based tools and orchestration solutions, IBM's services are designed to help clients protect IT systems, keep mission-critical applications running, and achieve fast and reliable recovery in the event of outages. For more information, visit: *ibm.biz/multicloud-resiliency*.

CONTENT PROVIDED BY: IBM

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
Chrysler Building
405 Lexington Avenue,
9th Floor
New York, NY 10174
+1 212 505 3030

SAN FRANCISCO
505 Montgomery Street,
Suite 1052
San Francisco, CA 94111
+1 212 505 3030

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 203 929 5700

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

451 Research®