

Securely Embracing the Cloud



451 Research

S&P Global
Market Intelligence

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

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 public cloud infrastructure has become a given for many organizations. The scalability, ubiquity and ease of consumption that cloud models offer can be a tremendous advantage. It ought to be a more secure environment for most. But the benefits can be tempered by a new set of risks that come along with the public cloud service model. While the security capabilities of most offerings are many, there is still the fundamental challenge presented by handing over responsibility to a service provider. It's true that there is the risk of malicious insiders in any environment, but that risk looms larger when direct control has been formally handed off. New technologies and services that have arrived under the banner of confidential computing now offer tools to address many of these issues, but they require a unified way of managing their capabilities and extending them back on-premises. Moving from confidential computing technologies to a confidential cloud approach, where those data protections are integrated transparently into infrastructure, offers a way for organizations to get the best of cloud – addressing the risks created by public cloud without changing the way they operate.

Key Findings

- Confidential cloud approaches can address both on- and off-premises risk.
- Confidential computing technology options exist in public cloud, but they need extensions beyond their memory-only coverage and management to reduce operational complexity.
- Protections that focus on securing data rather than infrastructure can simplify application security.
- Silicon-based security support, like secure enclaves, can support robust data security, but it needs to be extended to the full data lifecycle on-premises and in the cloud.
- Effective approaches to confidential cloud have to support a range of hardware and cloud options.
- Enterprise adoption hinges on reducing friction in augmentation of the existing IT ecosystem.

A New Set of Protections

It's no longer necessary to explain the benefits of external cloud to most enterprises. They get the value of scaling, flexibility and access to the latest technologies and services. For most, the physical security offered by cloud environments is better than they're able to achieve in their datacenters. This improved understanding has also brought up more sophisticated concerns about the use of cloud as a key part of their infrastructure. Working with key data assets in cloud adds an additional risk vector to their calculus. Expanding the exposure of that data to a new party, in the person of the cloud provider's staff and other tenants, requires a new set of protections. Shared infrastructure, with the potential for inadvertent or malicious exposure through privileged administrators or hostile code, can seem too risky. Some organizations have held off on cloud use, while others have worked to master native cloud protections or overlay their traditional, perimeter-based protections onto cloud with virtualized instances.

Figure 1: Top Issues with Hosted Cloud and Security Technologies Implemented

Q: What are the top potential issues with hosted cloud offerings (hosted private cloud, IaaS or PaaS)? Please select up to three.

Base: All respondents (n=199)

Q: Which of the following security technologies or capabilities – if any – have you implemented in your hosted cloud? Please select all that apply.

Base: Respondents who use hosted cloud solutions (n=198)

Source: 451 Research's Voice of the Enterprise: Information Security, Budgets & Outlook


These concerns are evident in the results of a recent 451 Research study. Data exposure tops the list, and other issues attached to data access and use make up a large percentage of the responses. Data security is clearly a large issue that is impeding the adoption and effective use of cloud services. This same study also asked respondents about their use of security controls that are available in hosted cloud services. While encryption was the second most common reported capability, greater priority was given to access controls in the aggregate. This indicates that much of the focus on securing cloud environments is tied to managing access to the various assets rather than protecting the most valuable one – the data that is being stored and used. Hardware security modules, which should be part of a data protection implementation, received the lowest score of any option presented. That's another indication that the focus of security strategies doesn't seem to be targeting what's truly valuable to the organization.

New Models of Data Protection Are Required for Effective Use of Cloud

There are many reasons that organizations view the task of implementing effective data-centric protections in cloud as difficult and expensive. Adapting existing, on-premises tools and techniques is certainly a complex task that has various challenges. Approaches that worked in a multi-layered, perimeter-based world can't extend to cover the risks to data in the cloud. Adopting cloud-hosted tools often requires new skills in implementation and operation. The cloud offers speed and scale, but those attributes can come at the cost of data security. To put the benefits of cloud-based infrastructure to work, organizations need to move beyond the false dichotomy of having to choose between having the speed of cloud or the reduced risk that their traditional data protections offer. If they are to thrive, they must get both.

The move to greater digitization of the modern business environment offers significant advantages, but it requires the integration of much more data in all aspects of an organization. A 451 Research Voice of the Enterprise (VotE) Data and Analytics study looked at what was holding organizations back on their digitization journey. After budget concerns, respondents cited data-related issues as four of the next five highest barriers. We compared these results with the results from organizations that see themselves as the most data-driven, in which data security and data privacy jumped up to the top four. This is an indication that organizations further along on their digital transformation journey see the need for data protection as a much greater part of their path to success. They also indicate that they're still hobbled by the tools at their disposal. By looking to new technologies for data protection, it doesn't have to be that way.

Figure 2: Barriers Faced in Attempting to Be More Data-Driven



Q: What are the biggest data integration challenges faced by your organization? Please select all that apply.

Base: Data management respondents (n=370)

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Management & Analytics

Getting to Data Protection by Default


Information security professionals are familiar with the idea of protecting data at rest, in transit and in use. Different tactics were employed for each: storage encryption worked for data at rest; network encryption worked for transit; secure application design worked for data in use. But they were all applied and managed independently. There is an additional concern in the full data lifecycle that has to be addressed, and it is more evident in cloud – the secure disposal of data. When organizations no longer have physical control of the systems that their data passes through, they need to address the exposure risk from reuse or decommissioning as well. All of these take significant work to operate, and they require awareness of where data will be located throughout each phase. The requirement to know and deal with this collection of special cases dramatically increases the complexity and cost of security operations, while still leaving gaps in protection.

Data protection has to move from being the exception to being the default – being part of the infrastructure. The traditional approaches treat data protection as a special case, where protections have to be applied to infrastructure in order to secure it. That's a result of early priorities to increase data sharing that didn't include the fundamental business requirement of data protection. It's also a side effect of the limitations of the current technologies that are used to secure data.

Existing approaches have individual aspects that are important to a comprehensive security strategy, but they often fall short of offering complete protection for the full lifecycle of data in the cloud. There are controls for access that can limit the exposure to sensitive data. There are controls baked into applications that deal with data use. And there are various forms of encryption that have been employed to protect data directly. The ways in which these are used today all have drawbacks that can limit their effectiveness. They wind up being layered on top of the supporting infrastructure and each other, adding more complexity than protection. Each establishes a perimeter that goes somewhere beyond the data that it's protecting, rather than being inseparable from the data itself.

Access controls are mandatory in any environment, but they have a number of limitations as effective data protections. Network-based protections, such as isolation and micro-segmentation, can prevent network connections to resources and services that provide sensitive data. Blocking access to a database or application function by identity or network address is a good start, but it can be complex to administer. As application development speeds increase, the speed with which network-based controls can be updated can be a limiting factor. These controls are also more subject to error because the teams requesting changes and those administering them are often separate. These controls are necessary, but not sufficient for data protection. The nature of access controls is such that when they fail, the data they are supposed to protect is still exposed.

Protections that are built into applications are another means of dealing with data security that can be effective. Data can be encrypted and only decrypted when in use within the application. This can manage risk for data at rest or in storage and in transit, but it's complicated to implement successfully. The management of encryption keys within the application has to be secured, and the development teams working on the application have to be sufficiently skilled to avoid implementation errors, which has historically been a weak point in application-based protections. Finding developers with the necessary cryptographic skills to build and maintain application-based protections can be challenging. A recent 451 Research VotE Data and Analytics study showed skills shortages and data security requirements as the top two reported barriers to integrating more data into their operations. Even if organizations could overcome these problems, cloud presents the additional challenge of having decrypted data available in memory on the systems where the applications execute. That risks exposure through maintenance by cloud administrators or an incident of compromise at the cloud provider.

Figure 3: Biggest Data Integration Challenges

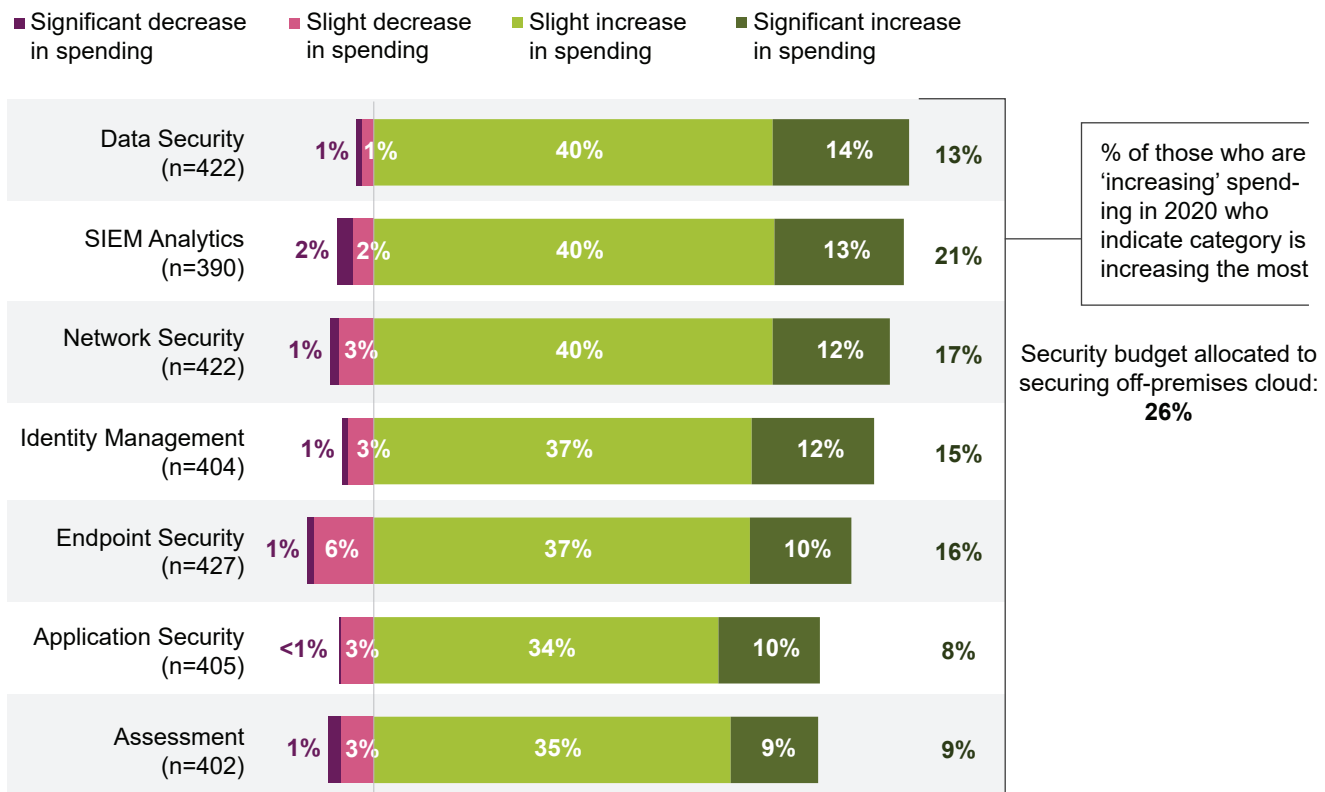
Q: What are the most significant barriers your organization faces in attempting to be more data-driven? Please select all that apply.

Base: All respondents, abbreviated fielding

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Management & Analytics

Data encryption can be managed externally, and cloud providers are expanding their capabilities to protect data in its various stages of use. The difficulty with using native data protection schemes is that there is the risk that data in different locations will become islands that aren't interoperable without considerable effort. Data that's in on-premises environments can't be easily integrated with that in the cloud, and differing cloud schemes add more complexity. Approaches like homomorphic encryption offer promise in bridging islands of data, but they have limitations that prevent them from being practical today. The computational work required makes processing impractically slow for most applications. Fully homomorphic approaches require complex designs. Even partially homomorphic techniques may require different encryption types for different use cases – dramatically expanding not only development work, but also storage requirements.

The problem is not that organizations aren't spending on data protection. Data security topped projected spending in 451 Research's VotE Information Security study. The problem is that those investments aren't being made in places where they can be the most effective. The complexity of existing environments consumes budgets as organizations wrestle with management and integration. They're spending on teams, as well as technology, as they attempt to make these technologies perform. As with many transitions in security, early attempts to make insufficient technologies do more than what they are capable of achieving results in high costs for those that devote resources, and inadequate protection for those that can't.

Figure 4: Expected Spending Changes for Vendor-Based Security Tools in 2020

Q: For each of the following vendor-based security tools, how will your organization's spending change in 2020 – if at all – compared to 2019?

Base: All respondents that will spend on vendor-based security tools in 2020

Source: 451 Research's Voice of the Enterprise: Information Security, Budgets & Outlook

Organizations have to transition to approaches that are practical. For most, application changes are too risk- and resource-intensive. The time and effort required to refactor an application to build in better data security capabilities can be daunting. This is especially true in light of the skills gaps that have been identified. Any workable approach will have to be part of the infrastructure supporting key applications if these challenges are to be overcome. This is even true in cloud-native applications, where insufficient data controls or architectural limitations could be just as vulnerable to exposure. They have to be protections that can be put to work in a way that's transparent to the applications themselves.


The Shift to the Confidential Cloud

One of the most significant changes in data security options is the shift to approaches that address full stack coverage for data. The work that started with confidential computing efforts has continued into more comprehensive approaches that can best be referred to as the confidential cloud – encompassing not just compute elements, but the full infrastructure. The roots of this shift have been driven by advances in the isolation capabilities provided by the next generation of processors and cloud computing environments that are now widely deployed. In the same way that virtualization extensions opened the door to better utilization, silicon support for greater isolation through secured enclaves now offers dramatic improvements in infrastructure security. Where protections like encryption had to be applied to the different realms in which data existed, these technologies can create secured execution spaces where data and applications can exist together, without additional modification, leveraging the capabilities or services that are native to the infrastructure.

When built properly, confidential cloud environments can protect the data that's being used within them by handling ingress and egress encryption. That allows data to be protected in ways that are independent of the application using it. The application and the data enter the protected space, and the work is done there. Data results can be handled in ways that meet the security criteria of the use case and can be independent of the protections used for the source data. It allows the protection perimeter to be drawn around the data itself, rather than at arbitrary technology boundaries beyond it.

Having a security anchor in the hardware that supports computing infrastructure enables a level of trust that can create opportunities for use cases that were difficult to achieve with traditional techniques. Enabling a common platform for the use of sensitive data that extends from on-premises environments to secure the use of external clouds and their confidential computing capabilities is just the starting point. Collaborative data analysis – where separate, private data sources are brought together, as in medical research – are possible because the source data isn't exposed during analysis, and only results data can be allowed to emerge from the secure environment.

Secure enclaves and confidential computing services are already available from several technology vendors and cloud service providers. They're offered in a number of models with varying levels of management support, which is something that their customers have been requesting. In a 451 Research VotE Cloud Hosting and Managed Services study, respondents rated securing their environments and achieving governance and compliance objectives as their two highest concerns. Both are areas where there was also a strong willingness to work with cloud providers through a managed service to accomplish these goals.

Figure 5: Challenges Implementing Cloud Technology and Likelihood of Working with MSP

Q: Which of the following challenges does your organization face as you implement cloud technology, services, platforms and environments?

Base: All respondents (n=364)

Q: Which are you likely to work with an MSP or professional services to address?

Base: Organizations facing challenges implementing cloud technology, services, platforms and environments (n=325)

Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Budgets & Outlook 2020

Organizations have worked with the shared security responsibility model that's offered by most cloud providers, but considerable effort is needed to bring that to levels that will provide deeper data security. These services can be a large step in making cloud environments more secure.

One of the limitations in the various offerings at this early stage of the market is that the different approaches each need data and applications to be handled in different ways. Refactoring of code may be required. Enclaves can protect data in memory, but storage protection has to be managed separately. For all of them, on-premises environments have to be managed separately. Potential users need to evaluate how well the different models support their needs or look for platform products that can span multiple environments with a common management system. The goal should be to build capabilities that can secure data wherever it needs to be used, extending confidentiality across infrastructure and geographies in a way that's agnostic to what's supporting that data. In the same way that organizations have simplified the management of encryption across on-premises and cloud infrastructure with bring-your-own-key efforts, they should be able to extend confidential computing capabilities with a bring-your-own-confidentiality approach. A tightly managed data perimeter has to be the end result.

Looking Ahead

As organizations continue on their journey to better utilization of cloud-based infrastructure, they need to expect to have securable data environments. Their focus needs to shift from the myriad protections that they have needed to deploy to self-securing data. Such a shift requires a change in thinking about the acceptability of disjointed security controls and the way in which data security has been approached. That change is embodied in the confidential cloud mindset. It's a shift from thinking that sensitive data requires special handling to thinking that sufficient protections exist in all situations for all data – that real data protection has to become ubiquitous, and that sophisticated data security services are just part of the infrastructure.

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.