

REPORT REPRINT

SASE, ZTNA and XDR: Three security trends catalyzed by the impact of 2020

AUGUST 19 2020

**By Scott Crawford, Garrett Bekker, Fernando Montenegro,
Aaron Sherrill, Eric Hanselman**

The COVID-19 pandemic could accelerate or hinder these initiatives, which embrace multiple technology segments. Will they be successful, or become mere positioning that could further frustrate the security buyer? That could depend on whether their components are complementary in practice.

THIS REPORT, LICENSED TO CISCO, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



451 Research®

Now a Part of

S&P Global Market Intelligence

Introduction

As we emerge from the first months of a very different world from the one prior to the RSA Conference, we find ourselves taking stock of the trends shaping technology and asking what's next. We recently had occasion to think about this question in detail as we updated our outlook for 2020 – not from the point of view of late 2019, when it was first published, but through a very different lens in the months since the World Health Organization declared COVID-19 a pandemic. Midway through Q3, we find ourselves taking stock once again as we prepare for the coming year, and we see some definite themes emerging in information security.

When it comes to the technologies involved, those themes can, on the surface, appear to be only marginally related: threat detection across endpoints, networks, clouds and applications; access control from any source, to any target; and the reshaping of security architecture driven not only by the ongoing move toward the cloud, but also by a profusion of new endpoints, including a plethora of operational and industrial devices – not to mention the millions of business users now working remotely for the indefinite future.

However, there is a common note sounding across all these themes. It has to do with a more holistic approach to enterprise security, as well as to enterprise IT. For years, we in information security have been talking about the need to consolidate a very fragmented collection of tools and technologies. Now, with larger trends driving change throughout IT, cybersecurity must adapt as well. In this report, we take a summary look at three of these – secure access service edge (SASE), zero trust network access (ZTNA) and XDR (where the 'X' is a placeholder for the intersections of techniques that contribute to threat detection and response, reflected by some through adoption of the term 'eXtended' for the placeholder) – and how the unanticipated events of 2020 are shaping their adoption.

451 TAKE

The coming world is being defined increasingly by technology delivered 'as a service,' with less dependence on product-oriented deployments and the fragmentation they introduce through components accumulated as capital investments in legacy environments over time. For security, this legacy has often meant that investments may rarely, if ever, be displaced – no one wants to be responsible for eliminating the one defense that is truly needed in the event of a certain type of attack. As-a-service providers have an opportunity to reshape this reality for security teams and consolidate functionality across multiple domains – and many of the trends we expect to follow going into 2021 will reflect this new landscape.

SASE: Redefining the boundaries of the enterprise

The rise of cloud computing has brought several capabilities within reach that would have been inaccessible to IT of the past, and in many cases more affordably. Although it comes at a cost, IT delivered as a service frees organizations from the long-standing burdens of maintenance and ownership. As a result, businesses depend much more on third-party providers for what they had previously had to deploy themselves, within their own environments.

Security technology has benefited from these advantages, as well, but there is one way in which security has also been challenged by such trends. Consider that many security tactics, such as access control, security monitoring, and network threat detection and prevention, have typically been deployed within the traditional enterprise environment in a way that matches each organization's idiosyncrasies. This is largely why remote access still depends to a high degree on VPNs. Apart from its ability to segment public from private traffic, a VPN also makes sure that enterprise traffic is subject to these controls, to assure the organization's security and policy priorities.

This, however, runs counter to the trend of embracing the cloud and third-party providers. In principle, any endpoint virtually anywhere should be able to access these services directly. Indeed, this ubiquity is one of the primary advantages of IT as a service. So how can organizations make sure that the same degree of control – over access policy, activity monitoring, threat prevention and confidentiality – can be applied to these new interactions that otherwise may not really need the intermediation of an enterprise network at all?

Enter the secure access service edge. As this collection of techniques might suggest, SASE (pronounced 'sassy') is not just one technology. It is a set of capabilities, many of which have long been familiar to the enterprise, made available as a service and accessible in principle from virtually anywhere.

This means not only security functionality, but the ability to deliver it as a service. Accessibility and connectivity are fundamental aspects of SASE providers, and their ability to deliver functionality at such breadth requires them to have a large enough collection of points of presence to eliminate the performance bottlenecks enterprises often face when setting up their own access architectures, such as the limitations of a traditional VPN or maintaining their own interconnection topologies.

Because SASE is such a collection of security techniques delivered as a service, with ubiquitous accessibility and connectivity as part of the package, it's a poster child for trends that represent this new movement toward more coordinated security functionality. Does this mean it's a technology segment? Not if it is made up of many established segments, where offerings within each are most directly comparable in features and functions – even if delivered by a single vendor. Is it a platform? Perhaps. One of the fundamental considerations of a platform is the functional interoperability of its components. This interoperability may be technological; the components integrate to greater or lesser degrees. This may not require its components to be delivered by a single provider; indeed, it may be beyond the scope of any individual provider to meet or exceed the capabilities of best-of-breed contenders in every SASE segment, at least currently. Multiple vendors can partner or otherwise coordinate on delivering a complementary set of capabilities, or it may be delivered as a managed service, where the provider handles the coordinated management of tools and practices to deliver a more holistic offering. Ideally, this coordination is largely transparent to the customer.


The recent wholesale move toward remote work has called out the need for such ubiquitous accessibility to security functionality more starkly, not least because of the need to extend these capabilities to millions of new remote workers and endpoints. Billions more endpoints are expected to follow, with the growth in enterprise operational and industrial devices expected to nearly double to almost 14 billion by 2024, according to 451 Research's Voice of the Enterprise: Internet of Things 2019 research. The need driven by the response to COVID-19 has put IT organizations under significant strain, as illustrated in our Voice of the Enterprise: Digital Pulse, Coronavirus Flash Surveys published in March and June.

REPORT REPRINT

Figure 1: COVID-19 Has Led to Significant Stress for IT Resources

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

Two-thirds of organizations are experiencing, or have experienced, an increased strain on IT resources as a result of the COVID-19 outbreak




Please indicate whether your organization has experienced each of the following as a result of the coronavirus (COVID-19) outbreak.
Base: All respondents

It's therefore hardly surprising to see information security increase the most between March and June as an area of expected spending due to COVID-19 among respondents to our Flash surveys.

Figure 2: Where Respondents Expect to Increase Technology Product/Service Spending Due to COVID-19

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



ZTNA: Toward evidence-based access control

SASE will play a role in delivering these capabilities, but one of the most fundamental – access control – goes beyond the definition of SASE. Access control requires the ability to apply policy to attempts to gain access to a given target. It incorporates concepts of identity; indeed, the term ‘identity and access management’ (IAM) defines the overall domain of information security, of which access control is a part. As such, access control may be a key functionality that SASE delivers, or helps to deliver, but it also transcends it. It may be more appropriate to say that access control intersects with SASE, but is neither defined nor bound by it. As part of IAM, access control is a field in its own right.

Access control has many moving parts, including functionality that processes an access attempt and determines whether to grant or deny access based on inputs and logic; an interface with the user or endpoint, to gather inputs such as credentials, device status and security posture, where policy may be enforced; and corresponding integration with access targets, to verify authorization, mediate policy-compliant access and protect against unauthorized access attempts.

Because of this complexity, developers often default to simple means to make an access determination: If the user or endpoint presents a recognized username and password, or originates from an enterprise network, it must be legit – not least because these approaches are often cheap and (relative to more sophisticated techniques) easy to implement.

As more than a few successful attacks have demonstrated, such techniques are not exactly the state of the art in access control. Credentials can be easily compromised, and even legitimate endpoints can be exploited by adversaries that know that network origin matters when seeking unauthorized access. Access control is one of the front lines of cyber defense, and organizations have strong motivation to improve against such tactics. This is turning them to more modern techniques to help make high-confidence access decisions.

The primary distinction between these modern techniques and those they seek to replace can be summarized as a matter of trust. If your policy is to grant access when presented with the right combination of username and password, then you’re trusting that whoever (or whatever) presented that combination is who they claim to be. It doesn’t matter if these credentials have been stolen; we trust that the attempt is legitimate because they have been presented. Similar logic applies to granting access based on network origin. Even if the endpoint has been compromised, it is on a network recognized as ‘trustworthy.’

More recent approaches seek instead to weigh the evidence – or rather, the combination of evidence that, taken together, makes a case for an access attempt being legitimate. Username and password? Check – if required at all, but not in isolation. Username, password and origin network? Okay, but what about additional factors? Can something else be supplied to verify the user, such as a one-time credential delivered out of band? What about the origin network? Is it a home office, from which access has been granted before without incident? Or is it a public place, perhaps originating from a browser on a computer, ATM or kiosk used by hundreds of people in a day? What about the device? What OS is it running? Is it jailbroken or rooted? What security features are enabled? Can biometric evidence help determine that the user is who they claim to be?

Hence the term ‘zero trust’ applied to such initiatives. The seeker must present sufficient evidence to make the case for an access decision – ‘evidence-based’ access control, if you will, rather than access predicated on simply trusting that simplistic conditions are enough. When applied to access sought across a network, the more definitive term zero trust network access (ZTNA) may be applied.

REPORT REPRINT


As the examples above illustrate, ZTNA may be more of a trend than a market segment, with components across multiple technology segments contributing to it. ZTNA, however, has had the opportunity to mature – relative to SASE, at least – to the point where specific guidance has been defined, such as the US National Institute of Standards and Technology’s recently released Special Publication 800-207 on Zero Trust Architecture or the Cloud Security Alliance’s Software-Defined Perimeter, which includes its own explicit reference architecture, with about 20 vendors currently providing a productized version in the market. These architectures, however, still span a number of functionalities. Endpoint systems can supply evidence about the configuration, software complement and health state of the endpoint, not unlike the Network Admission Control days of yore. Behavioral analytics can monitor activity to measure against determinations of ‘normal’ – to spot potentially malicious activity that would call for access restriction or refusal, or engage the gathering of even more evidence to increase confidence in access decisions. Technologies that play in the SASE realm can gather evidence about network activity or serve as policy enforcement points to mediate access decisions. Less explicitly or consistently implemented are fine-grain controls on access to specific target applications and resources – no small challenge, considering the range and variety of application architecture abounding today and yet to come.

Here again, the COVID-19 pandemic may be serving to accelerate adoption. Organizations may be reluctant to expand remote access significantly without raising the bar on confidence in access controls, potentially introducing increased complexity or new frustrations for users. ZTNA techniques that reduce these risks may therefore be the beneficiary of an immediate driver that may be motivating organizations to accelerate adoption that may have been in the works prior to the pandemic, but which has been prioritized to meet the long-term endurance expected of remote work according to our Coronavirus Flash Surveys.

Figure 3: Survey Respondents Expect Remote Work to Endure Long-Term or Permanently

Source: 451 Research’s Voice of the Enterprise: Digital Pulse, Coronavirus Flash Survey March and June 2020

Two-thirds of organizations now expect their expanded work-from-home policies to remain in effect long-term or permanently – up from 38% in March



Q: Do you expect any of the following measures that were taken in response to the coronavirus (COVID-19) outbreak to remain in place long term or permanently at your organization?
Base: Respondents whose organization took measures in response to COVID-19

XDR: A meeting of the ‘mines’

The shift in network security architecture driven by long-term trends such as IT as a service, as well as near-term factors like COVID-19, is having an impact on yet another trend shaping up in infosec: the bringing together of technologies and practices in threat detection and response.

Threat detection is not new to infosec. It's been part of both endpoint and network security for years. Signatures applied to malware and rules applied to network activity – both later complemented by machine learning – have played a role for as long as there has been endpoint antivirus or network firewalls and intrusion prevention. More recently, however, these technologies have matured. Endpoint threat detection regularly inspects deep inside the endpoint, detecting changes in configuration, software components or usage activity that indicate possible attacks or malicious activity. Network analytics, meanwhile, has advanced to leverage behavioral techniques that can discern malicious anomalies using more modern approaches.

Once again, we see a trend that spans technology segments in at least two domains: endpoint and network. They don't combine to make a new segment because the technologies themselves are markedly different. The techniques used to instrument endpoints for this deep granularity of detection reach into diverse areas such as configuring registries and system caches, while network behavioral analytics may rely on statistical analysis specific to network protocol dissection.

They can, however, be complementary in their use. When a security control such as an email security gateway or network perimeter defense signals a potential attack, analysts can turn to endpoint threat detection and response (EDR) technology to determine the extent of compromise. They can use EDR to determine if the endpoint has attempted to contact other hosts in efforts to spread an attack or assess the internal environment. They can also turn to network threat detection and response (NDR) to see if suspicious activity has been observed or to determine attempts at lateral movement. One might say they can ‘mine’ this information for detection ‘gold’ that can improve the timeliness and efficacy of response (if one were compelled to explain the somewhat contrived subtitle for this section of this report). They can further incorporate functionality to respond to malicious activity, such as automatically isolating endpoints or changing firewall or network segmentation rules to contain an attack.

Thus, while these technologies are not necessarily part of the same security technology segment, they are complementary – just as the ‘kill chain’ of an attack traverses multiple domains as an adversary seeks to penetrate its ultimate targets. The lack of coordination across these detective controls has contributed to the success of many past attacks. The combination of technologies into a more comprehensive initiative, such as EDR and NDR, along with telemetry and capabilities from areas such as email security, cloud and application technologies, and identity, has thus been embraced as an ideal to strive toward. Such integration is far from trivial, and countless hours and dollars have been spent on integration efforts. The rise of breach and attack simulation (BAS) systems offers a way for organizations to assess the effectiveness of their efforts and determine if the various parts are achieving the expected mitigations. They might even identify overlapping controls that could allow tool elimination, although that's a path many tread with caution. Could there be an easier way to get all this telemetry to somehow just work better together? This is the objective of XDR.

The promise of XDR is to make this integration seamless, provided you're using components from the same vendor (and that those components are well integrated to begin with) or a set of products made compatible through engineering partnerships between vendors (or through another approach, such as services that achieve effectively the same result). Once again, this is less of a new segment than an integration of existing technologies.

If the technologies of XDR are not a single segment, are they deployed as a platform? Again, it depends. As before, part of the definition of a platform is the functional interoperability of its components. This interoperability may be technological; the components may integrate to greater or lesser degrees. Also as before, this may not require its components to be delivered by a single provider; multiple vendors can partner or otherwise coordinate to deliver a complementary set of capabilities. With XDR, however, the inherent handing-off of some of this effort to an external entity has aligned nicely with the business model of managed security service providers. Managed threat detection and response offers comprehensive threat detection and response capabilities as a managed service, where the integration of process reflects the nature of an attack, its mitigation in response, and follow-up steps for containment, resolution, and hardening against future attacks. XDR fits the services opportunity like a glove; even the acronym allows for it.

So what's the impact of COVID-19 on XDR? Consider that many detective technologies – particularly in the network – have relied heavily, if not exclusively, on deployment within the traditional enterprise environment, deployed via network infrastructure the enterprise owns or operates within and between its own facilities. What does the shift toward IT as a service mean for network threat detection, and where can organizations regain telemetry that could be lost through such moves? Might the contributing technologies of SASE have a role to play in the evolution of XDR beyond the traditional enterprise?

These questions and more will motivate much of our research going forward, as trends (both long- and short-term) continue to reinforce each other in forging an overhaul of enterprise cybersecurity as we have known it.