Hypes and Trends in Digital Identities

Andre Priebe, CTO at iC Consult

Presented during the iC Consult IAM Pit Stop Series

Pit Stop #1: IGA

Strong digital identities are an important cornerstone of modern security strategies. But with the growing number of workforce, customer and partner identities, the management and the governance of the identity lifecycle is becoming increasingly difficult, and organizations are looking for technological innovations to reduce the administrative burden of Identity Governance and Administration. In his presentation at the first IAM Pit Stop Meeting, iC Consult's CTO Andre Priebe presented some of the most relevant trends from the recent Gartner Hype Cycle and their implications for IGA initiatives: He discussed the potential of Open ID connect and Document Centric Identity Proofing, dove deep into the topic of Identity Risk Management and closed out with exciting insights into new Micro Segmentation approaches. Brace yourselves for an exciting ride!

New technologies like to make bold promises, and it's not always easy to distinguish which of the emerging trends will really end up shaping our future. A great aide for this assessment is the Gartner Hype Cycle – an annual graphic representation in which the renowned analysts list and discuss the most important recent developments and their current maturity degree. During iC Consult's recent PitStop presentation, CTO Andre Priebe presented his own take on some of the upcoming identity-centric trends in the 2021 hype report.

Content

Identity and Access Management
Hype Cycle 2021

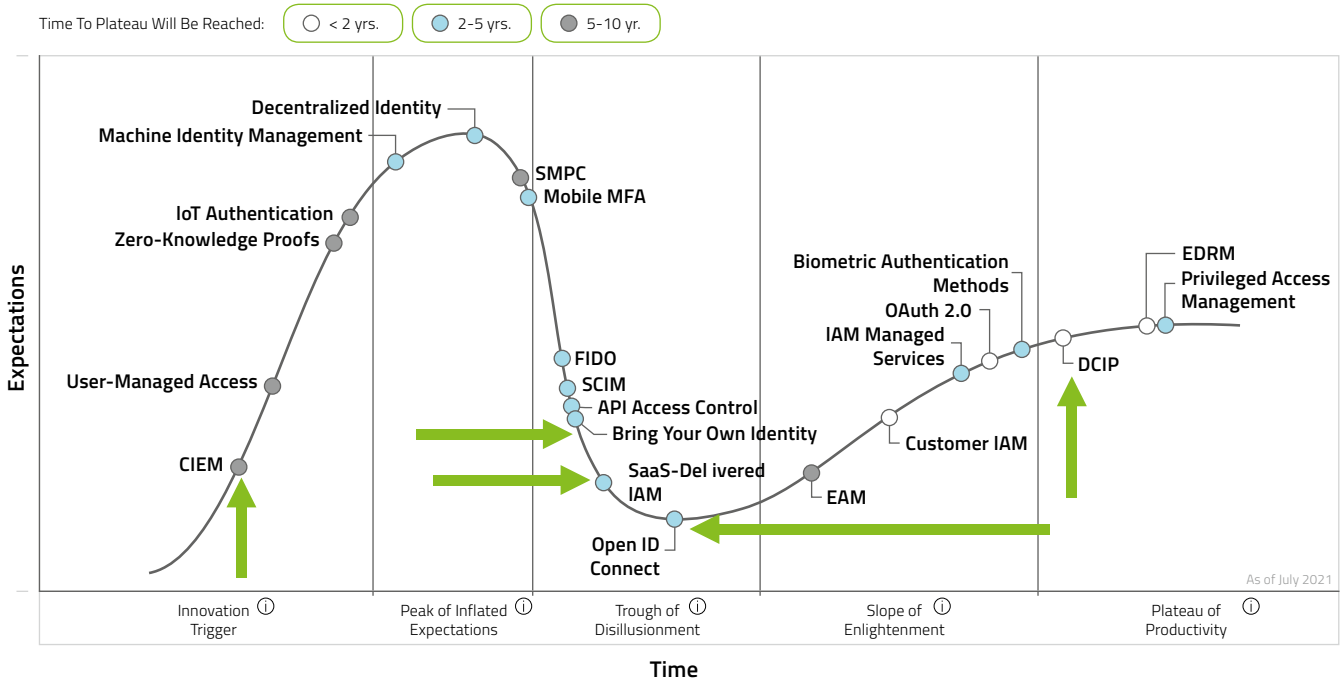

Hype 1: OpenID Connect

Why is OpenID Connect – which is not exactly a new technology – suddenly so relevant?

Identity proofing is not exactly a new topic. All of us have been implementing and leveraging the identity proofing technologies for many years. But recently, there has been a major paradigm change: We see that the responsibility to manage the proofing processes is rapidly shifting away from single business applications to a centralized identity access management system. This should come as no surprise: Typically, enterprises have to spend a lot of money to get an identity verified by an organization. Being able to use this information and this valuable data multiple times can save companies a lot of money. And what's even more important: It significantly improves usability.

Apart from that, there are other good reasons to consolidate identity proofing in a central spot: The first that comes to mind is regulations – but frankly, most of these have been out for years. Another reason is that identity fraud is increasing – and not just for customers but also with multiple attacks on the supply chain. A third factor is that the number of fully AI-based solutions for the verification of documented identity information is increasing. Everything is becoming more efficient, usability is improving and everybody's getting to know these solutions.

So, there are multiple good reasons why companies are looking into centralized IAM solutions, and a side effect of this is that a lot of things are happening around the OpenID protocol. This is still just in specification but already on a level where we recommend having a closer look.

The core idea is to enhance the information shared by OpenID Connect with a verified claims object which could provide additional information about which trust framework has been used. This will probably include:

- which rules have been used for the validation process – for example, specific money laundering laws or telecommunication laws;
- detailed information about the verification method that has been used:
  Was it physical, e.g., an in-person check at a point of sale or by an agent?
  Or was it unsupervised and remote – e.g., via an app scanning the document?
  Or was it by leveraging the electronic functions of the ID card?
- which organization was in control of the verification process; and
- which document has been used for that purpose, including its expiry date and a definition of the parameters you want to share.

OpenID Connect for Identity Assurance 1.0

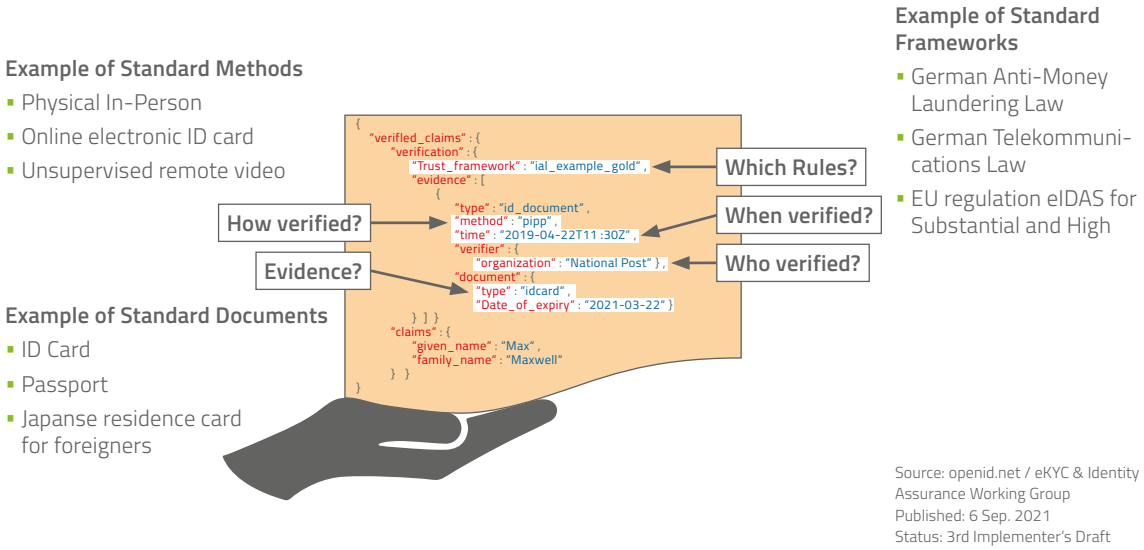

Source: openid.net / eKYC & Identity Assurance Working Group
Published: 6 Sep. 2021
Status: 3rd Implementer's Draft

This specification is already more or less final, and we are expecting only a few last adaptions. This is exciting news: OpenID Connect is promising to make it much more efficient to share identity-centric information within your organization. The big plus is that you will not have to take care of the solution design anymore and can simply leverage the existing specifications.

Note that this aspect is also becoming more important in B2B scenarios as well as in workforce scenarios where we cannot be 100% sure about identity anymore. This is especially important as a lot of people worldwide have been working remotely for two years now, and that's really driving the demand for strong identities.

Hype 2: Bring Your Own Identity & IAM as a Service

Organizations are often struggling when applying their IGA processes to external users. How can new technologies help us make those more efficient?

This is another big topic which has spawned several entries in the Gartner assessment. In this context, we ought to look at two major hypes: Bring Your Own Identity and IAM as a Service. In both of these realms, Microsoft is quickly becoming one of the more important players, and for good reasons, because they deliver identity as well as rich business functionality. A lot of our clients have already migrated to Microsoft 365, but their identities have not yet been perfectly integrated. Or maybe they have been integrated but they are also thinking about implementing additional interesting capabilities and features. How can we leverage that – together with our identity access management we already have in place – especially with regards to B2B applications.

So how does our typical IAM architecture look like today? Usually, an identity provider takes care of authentication, authorization, SSO, MFA and similar things. In addition, we have IGA, which is taking care of the identity lifecycle, the roles, processes to provide privileges and so on. And all of this is not only needed for our employees but also for all other digital identities accessing our IT systems. This includes partners and suppliers, external sales organizations or even customers. Combining all of these moving parts in one comprehensive solution is becoming more and more difficult.

Enter Microsoft 365 with its huge footprint. They have a very compelling offering, especially with regards to Active Directory. Why? Well, Microsoft has to provide and manage strong digital identities for many of their solutions already – including email accounts, shared documents and many other capabilities like the collaboration in Microsoft Teams. And they are doing a great job not just for internal teams but also when it comes to integrating external accounts like customers or suppliers.

One of the most exciting features in this context is "Guest Invite," which allows users to invite guests or team members based on policies to collaborate in a very, very efficient way. Now note that this is usually targeting the same people who are already in your partner or supplier identity repositories, your database directories, or other directories. So, when you are inviting guests into your working space, you will usually use their external email address – and that email address is often already known to Microsoft, because the invitees are also using Microsoft 365 in some way, shape or form. This is the foundation for a very innovative feature, which comes for free with your Microsoft tenant: the Federation. It's waiting there for you, ready to connect you at scale with thousands of enterprises, provided you are already operating with these companies.

B2B Identities and Microsoft Azure AD

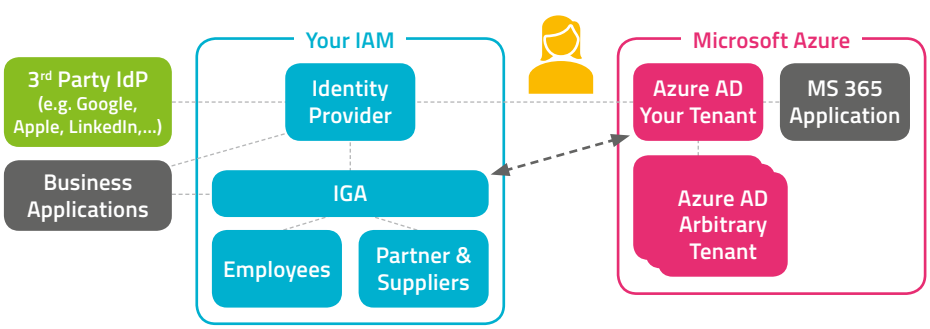

That said, there is still one major challenge, and that's the lack of a central authorization to define what access an external person has within your organization. We will have to find a way to integrate Microsoft Azure and IGA tools in order:

- to understand what kind of guests are out there
- to stay in control of the process required to invite guests
- to be able to proof identities for these guests; and
- to have a clear understanding of what kind of guests we want to allow and what rules we need

All of that boils down to one question: What kind of Federation partnership do we want to have? I really recommend our customers to discuss this with knowledgeable experts. Microsoft Azure is a very powerful tool. It offers a wide set of features and grants a lot of flexibility for your integrations. But it can also be a bit overwhelming at first, so you really need to address the topic. Otherwise, you will have no control whatsoever over what is happening with external identities on MS365.

Note that due to the huge amount of features and configuration options, you will rarely have full visibility of your identities right from the start. Also, Microsoft and Active Directory capabilities are developing rapidly, with new features being added all the time. So, the whole ecosystem is very complex but also very dynamic and that is, of course, a good thing – but not without challenges.

Hype 3: Zero Trust & Identity Risk Management

What is Zero Trust – and do we really need it?

Another huge topic, which has already made it safely into the late phases of the Gartner Hype Cycle is Zero Trust. The core idea of Zero Trust is very simple. Let's look at a scenario where one of your end users wants to access some resource via his notebook. In the past, he had to gain access to the network first. But once he was inside, he found himself in a large zone of implicit trust, where the resources were not necessarily unprotected but definitely not protected well enough to expose them to the Internet.

Zero Trust Architecture
High Level Overview

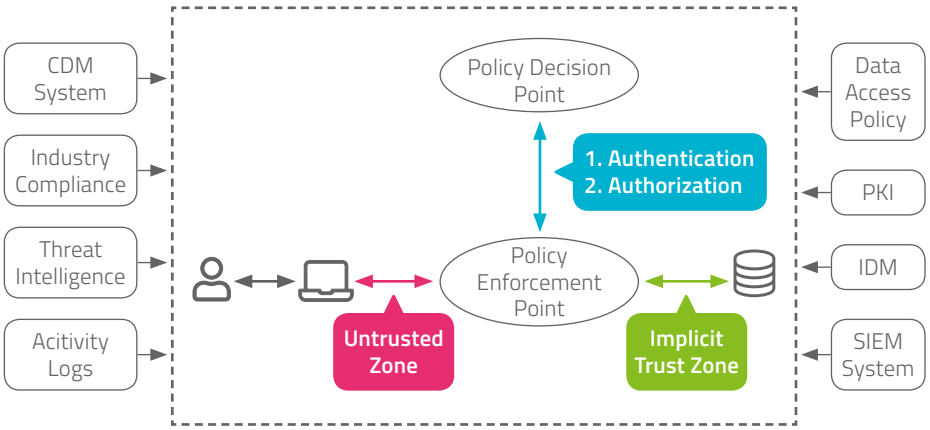

Source: NIST Special Publication 800-207, Aug. 2020

Now, the core idea of Zero Trust is to make this zone of implicit trust as small as possible and the untrusted zone as large as possible.

In the example above, this would mean having a policy enforcement point right before the resource the employee is accessing – basically: a spot where we will authenticate and authorize him. The next step will be to collect as much information about him as possible to allow for a robust authorization decision. Having a dedicated set of evaluation policies before every single resource will make it much harder for any attacker to move through your network or to jump from one compromised system to the next. That said, there are a couple of challenges when implementing this approach.

Challenge 1: Complex Policy Management
A big one is policy management, which is getting more and more sophisticated. Nowadays, most companies want to check, for example, what kind of mobile device a user is connecting with, before letting them access any information or to include dynamic threat intelligence in their decision making. All of this requires very smart frameworks. Another challenge is that you need to understand both sides and all implications of the decision-making process very well and need to comprehend both the technical and the business implications of your policies: If you start blocking non-compliant devices from accessing your network, this might heavily affect partners, suppliers and customers and is not a decision that should be made lightly. So, policy management can be a challenging task. It can cause ripples across multiple business processes and applications, and it depends strongly on your organization's risk appetite.

Challenge 2: Managing Policy Enforcement Points
The second key challenge is the policy enforcement point. While some vendors position themselves as full-scale providers for all things Zero Trust and claim to offer a fully integrated solution, that is not always the whole picture. The policy enforcement point is not a simple device but rather a complex concept: It needs access to all resources, has to speak and understand multiple protocols and different architectures and be able to interact with a vast array of different technologies – all while being tailored to protect that one specific resource. That's very hard.

Challenge 3: Lack of Visibility
But there's an even bigger issue: We can only protect what we see. We have to know it exists, because otherwise we won't even know that we have to put an enforcement point in front of it. This is something we have to focus on: How can we get full visibility of what is out there? How can we protect it? How will we know if there are changes or if new resources get added somewhere in the cloud?

This takes us straight to the next point when talking about your Zero Trust architectures: You have to make sure that these activities are really covering all your critical assets – even if that critical asset is an Excel sheet that should not even exist, with account and password information that could compromise hundreds of companies. But what can we do to ensure that our Zero Trust architecture achieves this goal?

The Solution: Comprehensive Identity Risk Management
To illustrate how to establish this kind of protection, let's have a look at the state of identity risk management today, why it is so important and which recent developments you can leverage to improve your organization's posture. We will look at the risk of account takeovers as an example.

- I think we all agree that in a large organization, account takeovers will occur. By default, the probability that the incident will happen is probably close to a hundred percent.
- Estimating the impact of such a takeover, on the other hand, is much harder. It depends on the tools and goals of the attacker and can vary a lot.

To reduce the risk potential for your company, there are two routes you can take:

- Some activities – e.g., introducing multi-factor authentication – will significantly decrease the likelihood of a successful takeover.
- Other activities can decrease the negative impact of a successful takeover. Zero Trust is a good example here: As long as you ensure that least privilege principles are applied in your network, this will significantly decrease the risk of lateral movement and limit the damage the attacker can cause.

Also, for both of these routes, there are several exciting new approaches that can help you mitigate risk even further. And while the following is not intended to be a comprehensive guide, it should help you understand how you can combine different procedures and technologies to minimize the risk of a truly devastating attack and tune the risk level so that it perfectly matches your organization's risk appetite.

Identity Risk Management

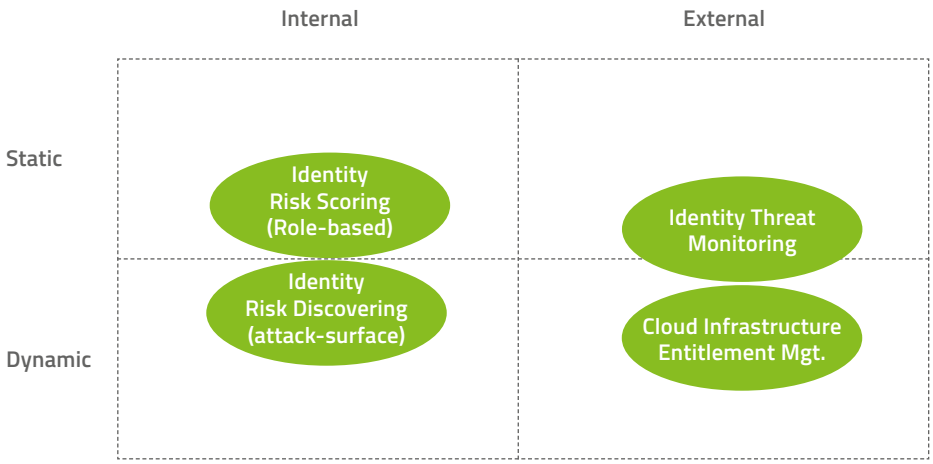

If we are looking to control our risks, our first step will usually be to figure out which users we have to pay most attention to. The traditional way to identify them is an **Identity Risk Scoring** (IRS) based on your IGA framework: To achieve this, you have to gain a robust understanding of the critical roles in your organization and the different roles every specific user holds. Based on that, it is easy to locate the 10% or 20% most critical users in your company and to implement additional layers of protection for them – e. g., by providing each of them with an impersonation resistant, hardware-based authenticator.

Another fairly new dynamic and sophisticated Risk Management approach is to analyze the tech surface from an attacker's perspective. There are some very smart tools out there today, which allow you to analyze the path an attacker could take when moving through your network. With these tools, you can look for hazardous files on notebooks, analyze connections, look where critical credentials are stored and get a very good idea of where your most critical systems and your privileged users are located. This approach is called **Identity Risk Discovery**, and it is a great starting point if you want to enforce an especially strict authentication for this set of accounts. A secondary benefit of this attack path analysis is that it will very likely make you aware of multiple high-risk accounts, which weren't on your original list. A good example are accounts which are used to authenticate to database accounts and usually have privileged access to this database; or local accounts with administrative privileges, which can escalate to gain additional administrative privileges.

A third important consideration in this space is the idea of **Identity Threat Monitoring** (ITM). Unlike traditional threat monitoring – which analyzes potential exploits an attacker might use to infiltrate your organization – ITM adds an interesting new twist: It focuses on the specific threats that target your identities, as identities are quickly becoming the new perimeter in a Zero Trust world. If you analyze the Threat Monitoring market, you will find more and more new solutions on the market focusing on the monitoring of identities, and those can be a very valuable tool to reduce the likelihood of an identity-based breach, e. g., via stolen or compromised passwords.

Our final hype entry in the Identity Risk Management space is **Cloud Infrastructure Entitlement Management**, or CIEM. If we have another look at the Gartner Hype Cycle, you can see it's a very new entry, and Gartner claims it will be five to ten years until this technology reaches its Plateau of Productivity. If you ask me, I think this estimate is a bit on the long side. CIEM will happen way faster. There are already some really sophisticated solutions available, which are very good at analyzing what resources are located in the cloud and which accounts and privileges can access them. Gaining this kind of visibility is becoming more and more important, and these tools are getting a lot of attention already. Another interesting aspect is that Microsoft is investing heavily into CIEM technology, and it will definitely become a key component for not only Microsoft Azure but also AWS and GCP resources. This really underscores the importance of the topic: A lot of enterprises are concerned with the lack of visibility and transparency they are experiencing right now. It has become so easy to add another fast service to your infrastructure – but obviously, whenever this happens, it introduces additional risks. CIEM promises to solve or at least alleviate this problem and will be an important cornerstone of future architectures.

And with that, let's return to the starting point of our analysis, the topic of identity risk management. If we combine the traditional and the new tools we discussed, we can significantly decrease the likelihood of a critical account takeover, because we know which accounts to focus on and can provide strong authentication and especially robust policies for these accounts. And we can also use these tools to decrease the impact of a successful attack.

And another aspect needs to be considered: As we have learned, some resources are not protected by our policy enforcement points so far – either because we are not aware of them or because we cannot cover them from a protocol perspective. This means we have to find different solutions for these vectors, and I will touch on some of these approaches next.

One such approach is not to put anything directly into the stream between the client and the resource, as the resource is a Windows file-share. Even if you are not using SharePoint, Box or similar tools, there are probably still file-shares with critical access in your organization, or a database which is used by users via Fat Clients, ODBC or JDBC. In these scenarios, we might only have the integrated Windows authentication or a legacy application which will not be replaced within the next year. One model to mitigate these kind of scenarios – which I like a lot, because it really increases the amount of protection we can provide via MFA – is an approach where we enhance the Active Directory capabilities with a plugin that is running directly on the domain controllers. The plugin kicks in whenever a session key is requested via ticket or when NTLM credentials are sent to the domain control from the resource. Once the first factor, the password, has successfully been verified, the MFA plugin will automatically enforce the challenge of a second factor, for example, the confirmation of a push notification. And only after this is successful, the session key is provided. This is exciting on many levels: We are working on a completely different protocol level here – one which doesn't care about the direct communication between the notebook and the file-share or the database but instead enhances the functionality of the domain controller. And while some users might find it hard to understand why they have to switch to their phone when accessing a shared file, there are some elegant solutions for this scenario, for example, sister apps which pop up on the desktop.

Hype 4: Micro Segmentation

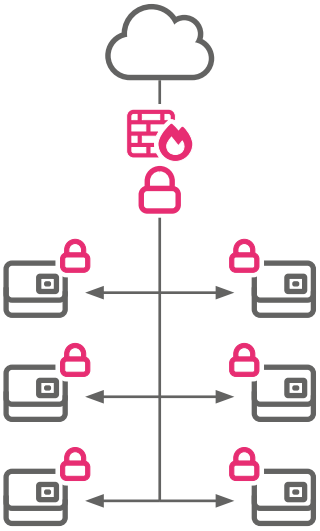

Isn't Segmentation a thing of the past?

And with that, I want to bring up one last topic that often gets overlooked when talking about Zero Trust architectures: micro segmentation. It is easy to explain why the topic doesn't always get the attention it deserves: Most organizations who implemented a Zero Trust architecture have already moved their resources from the old data center to innovative hybrid or cloud-based environments. Their applications run on Kubernetes or Docker, the traditional DMZ is a thing of the past and they don't believe firewalls will be sufficient to protect them there anyway. And that's why they don't really consider segmentation anymore. But they should! Because there are some very smart and innovative micro segmentation approaches.

Let me explain one of these models I find particularly interesting. As you know, a traditional firewall still works pretty much the same way it did 20 or 30 years ago: It is based on a rule set which regulates that it is possible to communicate from IP segment A to IP segment B; or that IP address X, port Y or protocol Z are allowed under certain predefined circumstances – and that's mostly it. You probably also know better than most that for a large IT organization, managing the firewall rules is quite a challenge, as the rule sets are expanding and it's becoming more and more difficult to understand which rule applies when or why it even exists at all. These questions come up all the time and the policies are really hard to understand for both the business users and the admins – especially in IPv6 environments.

And with that established, let's check what modern micro segmentation approaches can do for you. Let's assume we have a database system, for example, a logging system with a dashboard. Our user should be able to access that dashboard but not the database. And of course, the logging system should not access the database either – we just want the database to push its log events into the lock repository. So, basically

a very simple setup. Now, the traditional firewalling approach would be to control all these connections based on IP addresses and ports, but we don't want to do that. Instead, we will provide labels for each system: The database will be labelled database and maybe get a tag called 'Development,' because it's in the development stage (and not the production stage). Similarly, we have another system which will get the tag 'Logging' and a very wide area which we will tag 'Internet.'

Micro Segmentation
Alternative approach

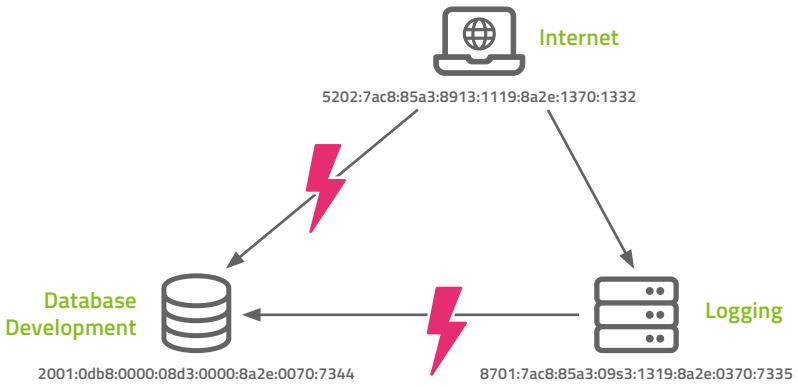

I think you understand the core idea already: We want to be able to formulate our policies in something fairly close to natural language: The new rules will allow a user to access the internet, the dashboard of the logging system or deny access to the database from the internet. These rules are very simple – especially when compared to traditional TCP/IP-based firewalls. But despite the simplicity, this approach solves most micro segmentation challenges in Zero Trust architectures. And what's even better: It is very dynamic and can be combined with strong identity-centric approaches easily. And it's not restricted to virtual machines or bare metal boards but provides us with a powerful tool for handling container environments, too.

The long-term goal for this approach will be to define the right number of tags – not too many, to keep the complexity low, but enough for the tasks at hand – and to use these to reduce the numbers of policies dramatically. Obviously, there are still some challenges: Once you start defining tags for every single system, the complexity and visibility of the rule set will become a major issue again and the tag surface is bound to explode. But the issue can be resolved, e.g., by having one tag for a whole project or a specific business application. This tag would then be assigned to all systems tied to that application and the team could proceed with the simple generic texts I have presented earlier. This model should pave the way for simple and concise rule sets that are easy to understand and review.

Now, to be honest, these solutions are still quite new, and no one really knows how they will develop over the next five or ten years. But while we don't have enough hands-on experience yet, the potential cannot be denied, and micro segmentation is definitely worth a closer look.

Conclusion

Robust, identity-centric Zero Trust security is currently in high demand. Log4shell has clearly shown that no IT systems are truly trustworthy. There are multiple vulnerabilities out there which we do not know yet and could affect a very large number of systems. Organizations should establish a comprehensive Identity Risk Management framework and proactively explore innovative trends and technologies like Zero Trust, CIEM and Identity Threat Monitoring to prevent dangerous and costly identity attacks. iC Consult can help you evaluate the different technologies and unlock their full potential for your organization.

About iC Consult

iC Consult is the world's leading independent consultancy, systems integrator, and managed services provider for Identity & Access Management with more than 800 employees worldwide.

We are committed to excellence and innovation, and with the best-in-class technology in the IAM space, we provide our customers with next-level cybersecurity solutions. Our service portfolio covers Managed Services for IAM including advisory, architecture, implementation, integration, support, and operations.

iC Consult is headquartered in Germany with offices in Switzerland, Austria, France, Belgium, Spain, Bulgaria, the UK, the U.S., Canada, India, and China. The world's largest brands trust in our expertise, to secure and manage their most valuable assets: their identities.

More information at www.ic-consult.com