Hypes and Trends in Access Management

Andre Priebe, CTO at iC Consult

Presented during the iC Consult IAM Pit Stop Series

Pit Stop #3: AM (Access Management)

The increasing reliance on a remote and hybrid workforce is forcing organizations to implement robust and compliant Access Management solutions. Zero Trust models are rapidly gaining traction, and security leaders are hard-pressed to offer solutions which combine strong security and a seamless user experience.

In his presentation at the third IAM Pit Stop Meeting, iC Consult's CTO Andre Priebe looked at the most relevant Access Management hypes and trends – and offered his perspective on an increasingly relevant and dynamic market. He spoke about exciting new developments around FIDO2, looked at OAuth 2.0 and Open ID connect, and discussed Risk Management approaches for modern Zero Trust environments. Brace for an exciting journey!

Content

New technologies tend 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:

Identity and Access Management
Hype Cycle 2021

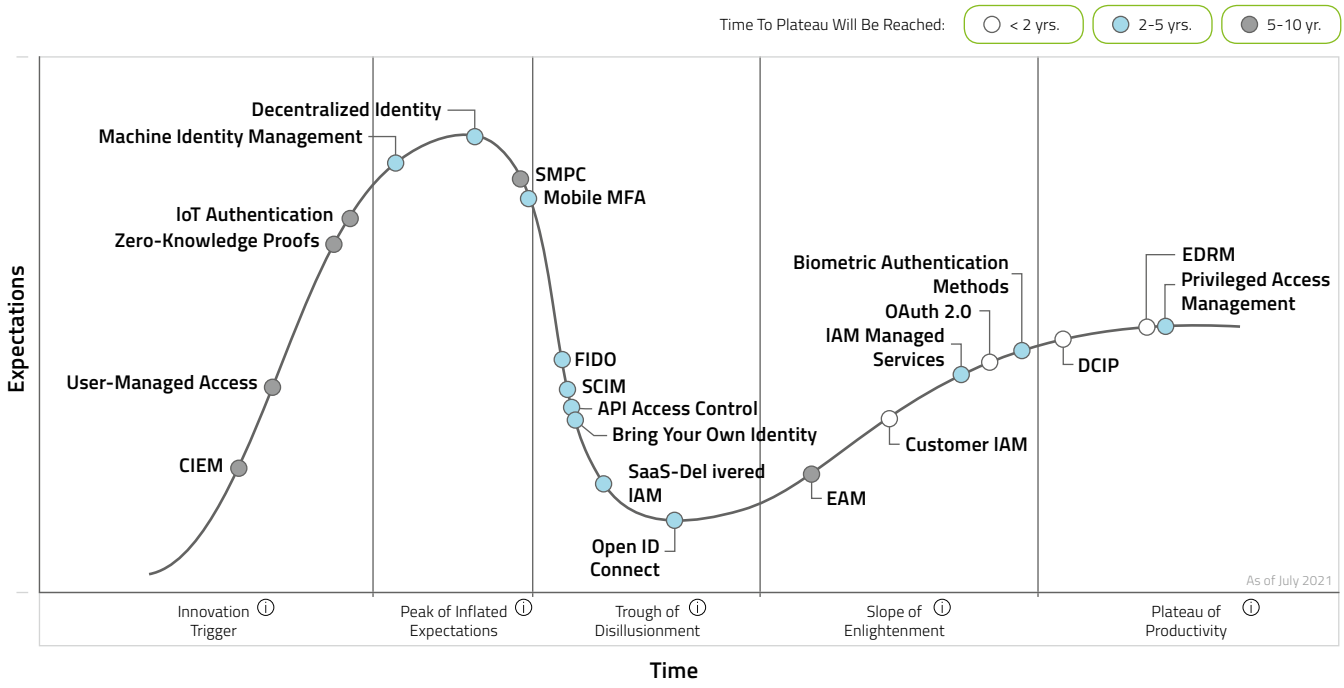

Hype 1: FIDO2

The two flavors of FIDO2, and why they might pave the way for a passwordless future

Let us start today's discussion with a look at authentication, because something very exciting has been announced in spring 2022 – and I would really like to share it with you and talk a bit about the impact I expect this will have. But before we dive in, let us look a little bit into theory, even if that is probably not new to you: When talking about authentication, we have three categories of authentication factors: The first is knowledge. The second is something we own – maybe a hardware token, maybe a smartphone, maybe something paper-based. The third is biometric. And depending on the category, there are multiple methods of authentication.

When implementing strong or Multi-Factor Authentication, we will always make sure that the different methods do not share any attack vectors. In other words, we will usually not ask our users for a password and a PIN, because both factors belong to the category of knowledge. Based on that method, we have a couple of concrete implementations out there: Some use single-factor methods, e.g., via password or a one-time code sent via email, or a generated One-Time Password. But we can also combine these methods with a PIN to achieve Multi-Factor Authentication, and there are also some methods which use a native Multi-Factor Authentication approach. A good example is a mobile authenticator app which combines push notifications with the biometrics or PIN protection of the local device. This will provide the security of strong two-factor authentication, but also the convenience of a single-step authentication. We can categorize these methods either based on the level of security they achieve or based on the level of usability they offer. And as we know, this rating will usually lead to heated discussions, even more so since it fundamentally depends on the implementation. For example – how would you rate the usability of a password? Most experts claim the user convenience is pretty good, as everyone knows how to use passwords. But on the other hand, no one is fond of passwords. And if you follow the best practices of password security and enforce both complex passwords and strict password rules, the usability becomes worse – so any ratings need to be taken with a grain of salt. And with that said, I want to talk about an exciting announcement related to FIDO2.

FIDO2 in a Nutshell

FIDO2 is an established authentication standard which has been published in 2019 and its adoption is growing at a very nice rate. Figure 2 shows how leading analysts are evaluating this technology – and as you can see, many of them believe that if you are thinking about introducing Multi-Factor Authentication into your company at all, FIDO2 is probably the way to go. And I would largely agree: Right now, there are still some legacy use cases FIDO2 is struggling with; but as far as future scenarios are concerned, I'm convinced FIDO2 will be the way to go.

If you are thinking about introducing Multi-Factor Authentication into your company at all, FIDO2 is probably the way to go.

So, it's definitely worth it to have a closer look at this technology. Basically, there are two distinct flavors of FIDO2, both of which combine something you have (your smartphone, your notebook) with biometrics (for convenience reasons) and a PIN (as a third option, in case you can't rely on biometrics): The first option leverages the TPM – the trusted platform which all of us already have on our smartphones or notebooks, the so-called companion devices. The great thing about this approach is that no additional hardware is needed, so the enrollment is very convenient and cost-efficient for the user. The second flavor uses dedicated hardware tokens which are integrated via USB or contactless NFC. This second method looks a bit more complex, but it does have one very important benefit: If you decide to use companion devices, you have to be aware of the fact that most users will keep their smartphones and notebooks for a maximum of two to three years. Then, they will get a new device – but their credentials will be tied to the old phone, and right now, that's still a major problem. Because if we assume that your average user will have two to three devices during his time in your company, that's how often you will have to enroll a new companion device for him. That's obviously neither ideal from the user's perspective, nor from the company's perspective, and might be one of the reasons why many organizations hesitate to lean into FIDO2 right now. There are some early adopters out there, but FIDO2 is definitely not the standard solution the analysts are making it out to be yet. For most organizations, it's just an additional option or alternative.

FIDO2 is definitely not the standard solution the analysts are making it out to be yet.

Multi-Factor Authentication and Password-less
Security vs. Usability

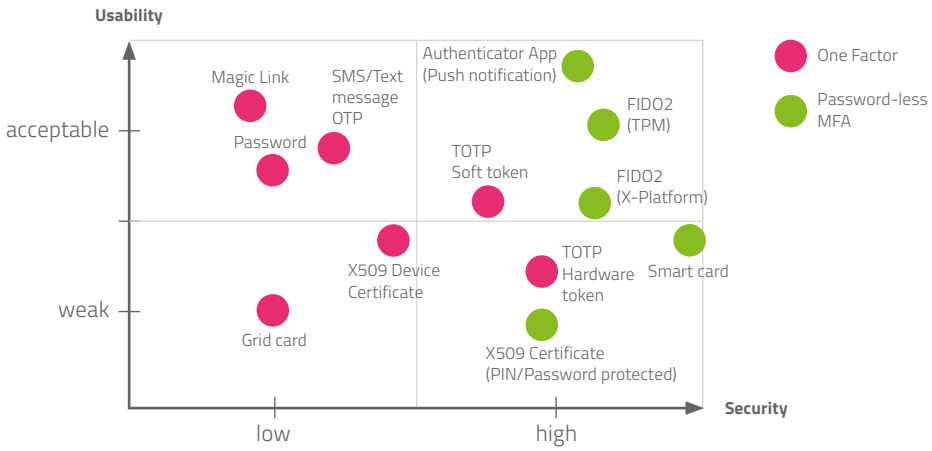

The other reason why it hasn't been adopted quite as much is that even for the same application and the same use case, the FIDO2 UI will look completely different on different browsers and different operating systems. It will behave differently, and the wording will be completely different, too. And what's even worse, you have no influence on the product and the UI – the only ones who can control it are the vendors who develop the operating system or the browser, and that's a major challenge: In order to provide a frictionless solution for your end users, you have to figure out which browser version, which operating system and which TPM technology they are using.

So, from the support perspective, there are still quite a few challenges, and you need to have a robust emergency process in place in case anything goes wrong.

The big news, and what they mean for us

With all the major criticism out of the way, let's finally talk about the big, new, exciting story: In May 2022, Apple, Google, and Microsoft published a joint announcement about the FIDO2 standard, and committed to expanding their support for the standard, in order to accelerate the availability of passwordless sign-ins. They agreed that they will introduce the capability to synchronize the FIDO2 credentials on a phone or notebook – so we are talking about the FIDO2 flavor without additional USB hardware – via their key chain in the cloud capabilities across all of a user's devices.

So, for one, this will cover the scenario we mentioned earlier: If a user's smartphone is broken, they can simply use their backup and migrate their credentials to their new device. From the usability perspective, this will be a major game changer – even though we have to be crystal clear about the fact that it will weaken security, as the authentication is not bound to the hardware device anymore.

This is the best chance we've ever had to get rid of the password altogether.

That said, I think this is the best chance we've ever had to get rid of the password altogether – at least of the passwords we are asking from the users; Apple, Google, and Microsoft will probably still continue to leverage passwords in some way, shape or form. And there's also the exciting option to use the smartphone – as long as it is nearby – to log in on the notebook. This will allow us to use strong authentication for the log-in without installing dedicated apps, leveraging only the built-in capabilities of the operating systems, really. That's a strong value proposition, and should drastically improve the user

experience, too. Of course, it still remains to be seen how fast this will become available in the market. Most likely, it's not going to happen this year, but hopefully next year.

And with that, let's continue our discussion about authentication and identity assurance, and about trust.

Hype 2: OpenID Connect

OpenID Connect for Identity Assurance 1.0 promises to lift the protocol to a new level – but the new features come with a note of caution

It is very important that you provide your organization with a framework that helps you understand which kind of trust is needed for a specific business application – based on risk appetite, use case, and regulations, of course. From my perspective, this is a prerequisite for any holistic Zero Trust approach. The x-axis in the figure below illustrates how we can define a precise level of identity assurance before we grant certain access rights – and how we can adapt this level for different use cases and business applications. In the context of Zero Trust, this means that we need to look beyond authentication and leverage any additional information available (about the location, the device, the user behavior etc.) before we make our access decision.

Authentication Assurance vs. Identity Assurance

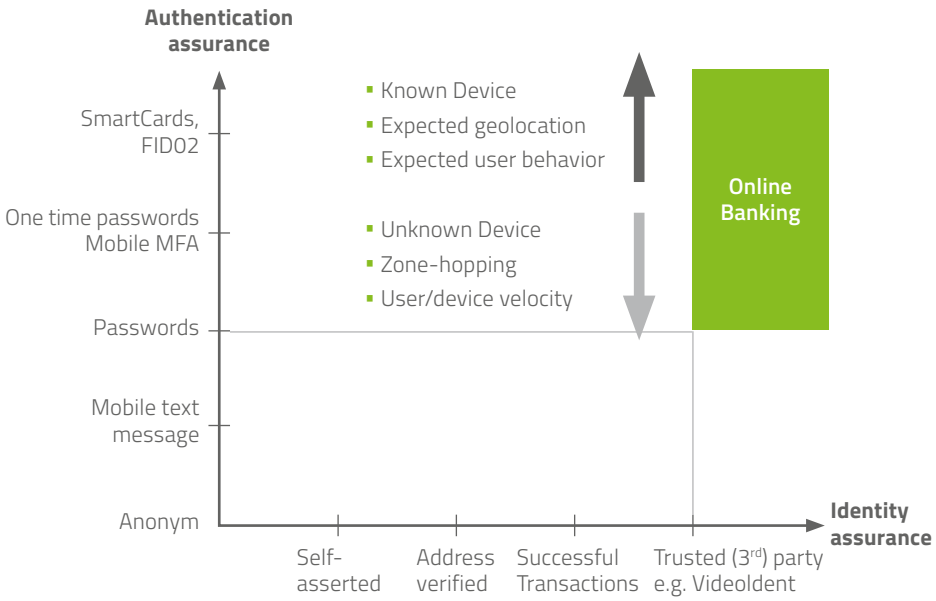

© 2022 iC Consult Group GmbH

And that brings us to very important new activities around OpenID Connect, which has been supporting this high level of authentication assurance from the very beginning, as a core part of the protocol. But what was completely missing was the aspect of identity assurance: OpenID could deliver the name, the email address, or the phone number of the user, but there was no standardized way of providing the accompanying identity assurance information.

This has recently changed with OpenID Connect for Identity Assurance 1.0, a new enhancement of the core specification, which is now supporting all the standard claims needed to provide information about users. In the past this has only been relevant when we were looking at use cases in Customer Identity Access Management. But nowadays, many companies have never seen some of their employees in person – and in these fully remote workforce and B2B scenarios, identity assurance is becoming more and more important, too, and OpenID Connect for Identity Assurance 1.0 is a protocol which provides you with the capabilities to deliver this verified information.

That said, OAuth 2.0 and OpenID Connect come with some highly relevant challenges that we should have a closer look at.

OpenID Connect was designed as an easy-to-understand protocol, but is becoming more complex as new standards are introduced.

Challenge 1: Token Design

If you recall the Hype Cycle we looked at, OpenID Connect is at the very bottom of the cycle right now. The reason is that it was designed as a very simple, easy-to-understand protocol – but with the addition of new standards like the ones we just mentioned, it's getting progressively more complex. And the original simplicity has led to another challenge: Many organizations have implemented OpenID Connect and OAuth 2.0 because they were looking for a simple solution, and so they didn't put too much thought into aspects like the token design. But now, after integrating hundreds or even thousands of applications into their Access Management, they suddenly experience performance problems or new use cases, like distributing their identity solution to new regions all over the world, and they find that the token design they chose a long time ago does simply not support this kind of scenario. Obviously, this will result in some disappointment.

Challenge 2: Confusing Roles and Scopes

Another major issue is the confusion of roles and scopes, two different dimensions when it comes to authorization: Roles are focusing mostly on the users and their rights, while scopes are looking at the application and what it is allowed to do on behalf of a user – which is a completely different perspective. Unfortunately, there are many situations where scopes are suddenly treated as roles because they are easier to integrate at certain enforcement points – and whenever that happens, it will add additional complexity in the long run.

Challenge 3: Functional Limitations

The third challenge are functional limitations, driven, for example, by browser privacy behavior. It's easy to imagine a situation where different browsers are developed into different directions, and this could easily disrupt the functionality of your OpenID Connect authentication.

Challenge 4: Upcoming Standards

The fourth challenge relates to the upcoming new authentication standards, some of which we've already mentioned. As discussed, these new standards are really driving the complexity to a new level, and this will mean that OAuth 2.0 and OpenID Connect will not be a reliable solution anymore. Just imagine that one part of your product will support one specific version of OAuth and another part of your product will support a different OAuth specification – then there's a very real possibility that these two parts will not be able to communicate with each other as they rely on different enhancements or interpretations of the framework.

And with that, let's have a look at Zero Trust.

Hype 3: Zero Trust on the Rise

How a diligent risk management jumpstarts your Zero Trust migration

The core idea of Zero Trust is to make the zone of implicit trust as small as possible and to have the enforcement points as close to each critical resource as possible. Now, let's have a look at some of the more exciting aspects and challenges of Zero Trust – and what could be a more exciting starting point than Log4Shell?

I think Log4Shell has taught us a lot about what it means to move the enforcement points closer to each resource. Since so many of our IT systems proved vulnerable to providing remote shell functionality, we had to ensure there was a policy enforcement point for every potentially affected resource. This enforcement point would be in charge of all policy decisions, including authentication and authorization of the end user, and following a set of activities to decide what kind of resource should be accessed. Since there are usually no policies to apply for these kinds of decisions, the subject itself and its roles become increasingly important – and this is an element that we will not get rid of that easily. Of course, there are situations in which you might implement externalized authorization, payback approaches or the like, but typically, roles will be involved in this process – to get privileged access by requesting it, for business approvals, and so on.

Zero-Trust Architecture
High Level Overview

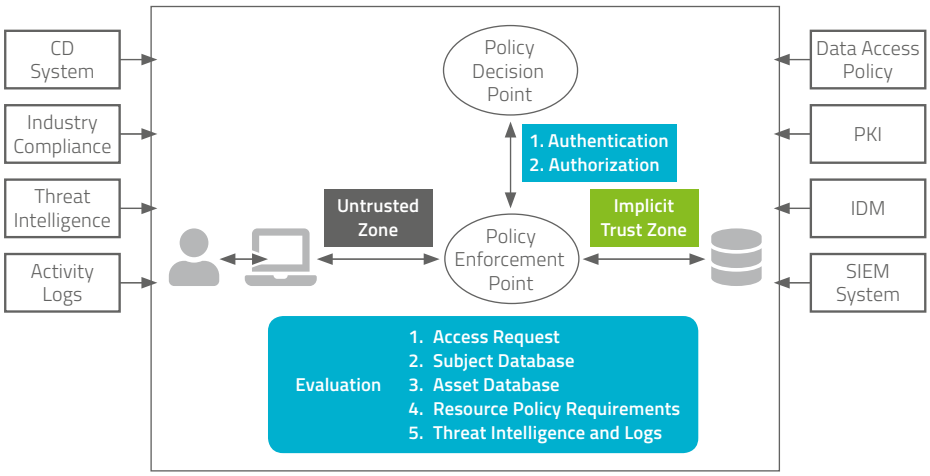

Another important evaluation tool in Zero Trust models is a comprehensive asset database. We often monitor the devices used for accessing resources – to see, for example, if the device is part of a device management solution – and leverage this information to evaluate compliance. This might include checking if the operating system is up-to-date or verifying if a smartphone has been jailbroken – and depending on the results, we will or won't allow it to access a critical resource.

Then, we evaluate access requests based on our resource policy requirements, and that's a particularly tricky component – because it requires a robust classification of enterprise resources, which remains quite challenging for many organizations. It also

requires a good understanding of your business cases and your specific requirements: On the one hand, it's great if your classification correctly defines a resource as confidential and your policy stops any unmanaged devices from accessing these confidential data. But what if you are audited externally, and the auditor is prevented from accessing highly relevant data by this very policy? Of course, you could hand auditors dedicated devices for their activities – isn't Zero Trust supposed to support all the different use cases in your organization? So, we should be able to say: If we have users from a specific group, then we will use different rules for their authentication and authorization – so that our access decision is not dependent on the device status, but some other rule – maybe some very secure, phishing-resistant authentication method used for privileged users.

Zero Trust is not a product you could buy, deploy, and be all set. Zero Trust is a journey.

And with that, I want to highlight one important aspect of Zero Trust projects: It's really hard to decide where to start – because unfortunately, Zero Trust is not a product you could buy, deploy, and be all set. Zero Trust is a journey. It requires a lot of thoughtful integration steps into your organization's existing IT systems, and you will have to grow your trust architecture step-by-step by adding more and more resources and information until you arrive at a coherent picture.

So, where should you start? Obviously, if you had unlimited resources, it would be easy – but that's not the case for most of us. So, let me share an approach we successfully leveraged for multiple clients already, and which tends to yield very good results in the early stages of the Zero Trust projects:

I would suggest focusing on some of the more relevant business risks you can prevent by implementing Zero Trust. The first use case is a compromised account in a business-critical system in a AAA security zone; the second is a compromised account in a training app provided by a SaaS provider. Now, I think we can all agree that the business risk of these two scenarios is very different.

That said, one thing is for sure: If your company is large and following fixed email and account name patterns, one of your accounts will be compromised sooner or later. Obviously, we want to reduce this risk, which is why most of us introduce measures like Multi-Factor Authentication. But even that will probably not be good enough to decrease the likelihood to "unlikely". As we learned a couple of months ago, even in scenarios where Multi-Factor Authentication with a mobile app and push notifications has been enrolled and enforced, particularly persistent attackers can sometimes trick users into granting them access. That said, a Zero Trust architecture can definitely reduce the impact of a breach.

So, in our Zero Trust world, we are always assuming that a breach is happening and looking for ways to mitigate its impact. This is where other components come into play: For example, we could use micro segmentation to make sure that the compromised system cannot access any other valuable system. Or we could apply the paradigm of least privilege to make sure that the compromised account will not be allowed to do anything exciting. And from there, you could even go a step further and implement additional security measures – maybe smartcards or hardware tokes – for the VIPs among your users. As for how to identify these valuable users, we'll get to that in a second.

But first, what about the other use case? The other use case is equally likely to happen. But we will probably all agree that due to the minor impact of this scenario,

the measures we have already put into place – basically: Zero Trust with Multi-Factor Authentication and least privilege access – will probably be good enough already, and there is no need to spend more money.

So, how can we identify the users who require special protection: One method is to focus on the internal perspective, and the very static side, and assess the risk based on privileges and roles the end user has. This will help us understand which end users have the highest privileges and to focus our activities on those accounts. That's the old approach. The new approach – which is much more closely related to Access Management than to IGA – is analyzing the actual attack surface of your accounts, based on the usage of your IT systems. Let me give you a concrete example: The solutions in this segment are scanning, for example, your notebooks or servers to find out which accounts are active on these systems. They analyze which credentials on the system could be used by an attacker to get access to further accounts and which accounts could easily escalate their privileges to domain administrator rights. This kind of analysis will give you a very good understanding of the system but also of the accounts – and thus help you protect your accounts and your resources more effectively.

The new approach is to analyze the actual attack surface of your accounts, based on the usage of your IT systems.

The next big thing will be similar tools running completely on an external site – solutions which are not running on any of your applications or infrastructures but are still scanning all the data available out there. They still scan the endpoints we expose but they will also look much further than that. For example, for any compromised accounts of yours that are actively being traded on the dark net – not because your Identity and Access Management solution was hacked, but because your employees are probably using their corporate email address to access a multitude of cloud services, and if any of these are breached, your credentials will be available on the dark web, too – and they will be bought by attackers trying to gain access to your organization. But seeing these credentials floating around out there is valuable information for us, too, as it points us directly to the accounts we have to protect. So, if you want to know more about the tools and products available in this category, please reach out to us.

Hype 4: Bring Your Own Identity

BYOI has been a CIAM mainstay for quite some time – but it does have exciting uses in workforce scenarios, too

And with that, let's now talk about Bring Your Own Identity – a trend we know very well from CIAM, and which has one very important use case in Workforce Identity and Access Management, too.

Companies like Google or Microsoft usually only provide the information that a usage has been authenticated by their organization.

So, what is BYOI: If you look at companies like Google, or especially Microsoft, they will typically provide us with the information that the use has been authenticated somehow by their organization and usually share self-asserted or managed data. That's not ideal, because we usually don't know where the data came from and how reliable they are.

That said, leveraging Microsoft, especially, is extremely beneficial – because, if a user brings their account from their home organization, it is a very safe bet that this account is well maintained. The best example for this is the deprovisioning process: Let's be honest – how long will it take you to learn that a contractor working for you is no longer with the company he used to work for? If you have federation in place, this will

be much easier: Microsoft Azure Active Directory will allow you to establish implicit relationships with other companies right out of the box, and it's very easy to use, too. So, BYOI with Microsoft is something you should keep in mind, both with regards to security and usability.

Bring Your Own Identity

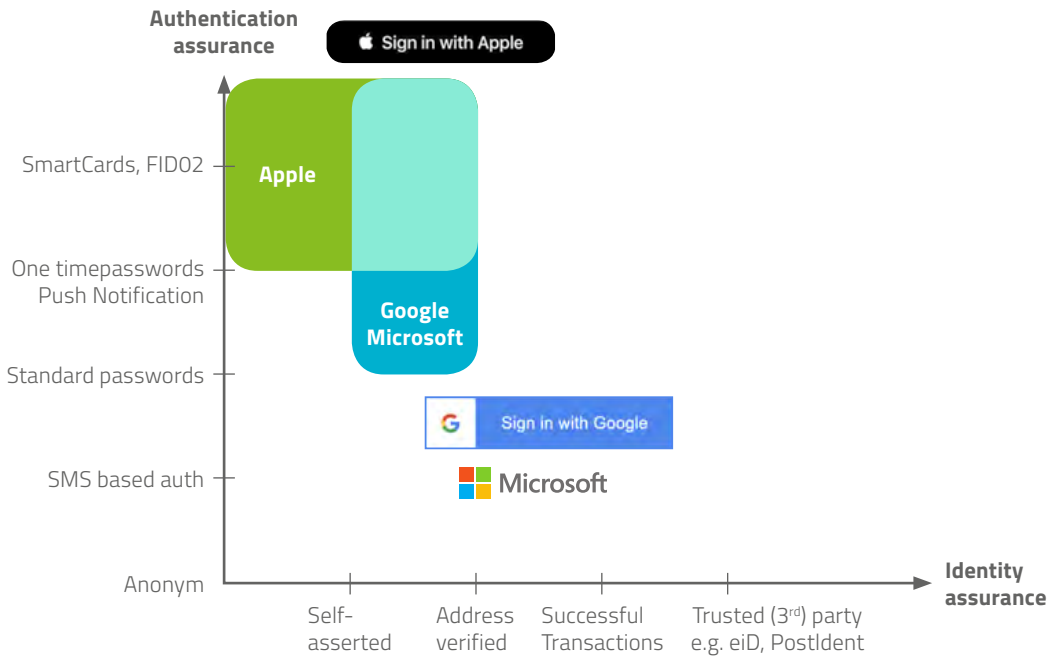

Apple, on the other hand, is used mostly in customer and consumer scenarios. That said, don't count them out completely, especially in BYOD environments. What they bring to the table is very reliable authentication, as they always enforce Multi-Factor Authentication – and that could be a big plus in some specific scenarios, for example when providing a job portal for external users where strong authentication is especially relevant. That said, there are some caveats when it comes to Apple. The biggest one is that you will not always get the user's real email address from Apple – sometimes, they will only provide a self-generated pseudo address and make it much harder to reach this contact.

And there are some more, vendor-neutral shortcomings when it comes to third-party login: For example, users might be using Microsoft AD in their home organization and log in with simple, password-based credentials there. Then, you onboard them and start enforcing MFA for your applications – which will probably be hard to understand for them and will also not provide the best user experience. Unfortunately, this kind of scenario is really hard to avoid at the moment – but we hope that Microsoft will provide the additional information we need to fix it in future.

Conclusion

Managing authentication and identities for remote workforces and global business ecosystems is becoming more and more challenging – and organizations need powerful Identity and Access Management solutions to secure and enable their worldwide user bases. Technological innovations like Zero Trust and passwordless authentication promise to change the way we work – but some of the most exciting developments are still in their early maturity phases, and the integration journey should not be taken lightly. iC Consult is excited to help you evaluate the different technologies and realize their full potential.

About iC Consult

The iC Consult Group, headquartered in Munich, Germany, is the world's leading independent advisory, systems integrator, and services provider for Identity & Access Management (IAM). The service portfolio covers advisory, architecture, design, implementation, and integration to IAM managed services and identity as a service offerings. The company's more than 600 employees have successfully delivered over 3,000 projects and managed services for IAM. The iC Consult Group has offices in Germany, Switzerland, Austria, Spain, Bulgaria, the UK, the U.S., Canada, and China.

More information at www.ic-consult.com