Privileged Access Management (PAM)

Andre Priebe, CTO at iC Consult

Presented during the iC Consult IAM Pit Stop Series

Pit Stop #4: PAM

The increasing reliance on Cloud infrastructures and DevOps environments 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 fourth IAM Pit Stop Meeting, iC Consult's CTO Andre Priebe looked at the most relevant Privileged Access Management hypes and trends – and offered his perspective on an increasingly relevant and dynamic market. He spoke about exciting new developments around CIEM, looked at PAM in a DevOps setting, and discussed Risk Management approaches for modern Zero Trust environments. Brace for an exciting journey!

Content

Hype 1: Cloud Infrastructure Entitlement Management (CIEM) 2
Hype 2: Privileged Access Management (PAM) 3
Conclusion 14
About iC Consult 14

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

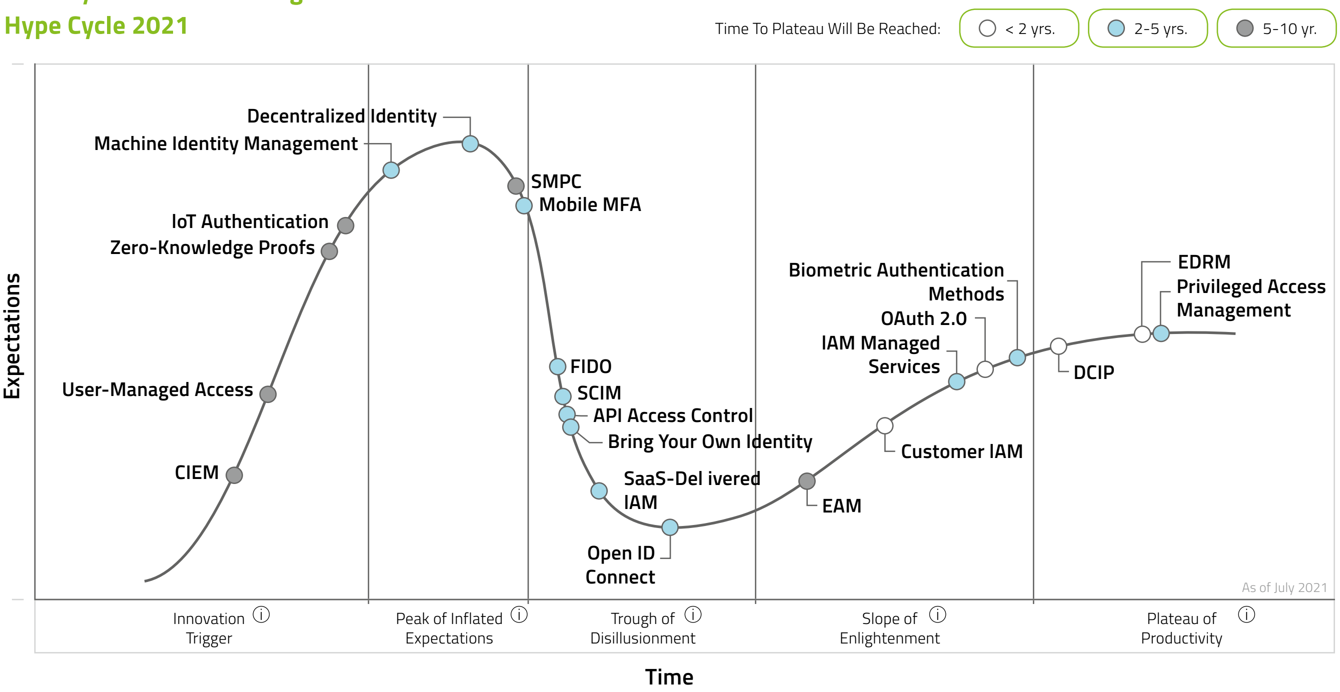

Hype 1: Cloud Infrastructure Entitlement Management (CIEM)

CIEM is still in its infancy but promises transparency of resources and entitlement in complex Cloud infrastructures

A lot of the resources we are protecting are part of a Public Cloud infrastructure service like Microsoft Azure, AWS and GCP. Do you have an idea how many different services are provided by Microsoft Azure or AWS? More than 200. 200 services for which you not only need an identity – credentials – to access them but also an idea of the kind of privileges required to administrate them. Which policies within your enterprise would apply to managing these services? Are they business-critical? Is there one role or are there 50 roles? How are the privileges split up within any given service? As you can imagine, having transparency of all these services provided by the Public Cloud infrastructure provider is very challenging – as is the integration into your Privileged Access Management (PAM) system.

This is where Cloud Infrastructure Entitlement Management (CIEM) comes in. CIEM explicitly focuses on these Cloud resources, providing an understanding of the services configured for your organization and the identities having access to these services from an administrative perspective – privileges, policies, resources but also the activities associated with these resources. CIEM solutions typically chart a permission graph in some way, shape or form, visualizing the critical privileges within your PAM system for you to sift through and tend to.

Cloud Infrastructure Entitlement Management

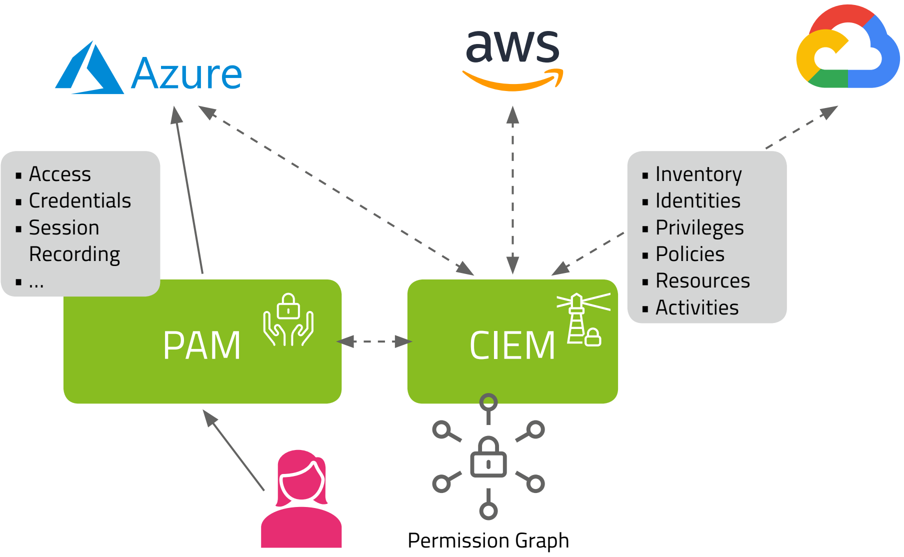

The Promise: Complete Transparency
The upside here, of course, is that the solution provider will do the heavy lifting for you, making sure that newly introduced or enhanced services are analyzed and then added to the permission graph as well. Otherwise, you never have a clear understanding of the percentage of identities, resources and privileges that is in fact covered by your PAM solution and the one that is completely invisible to you and your organization. That is what makes CIEM such an exciting and dynamic field. As of right now, there are a couple of small and larger vendors out there offering a sizeable set of functionalities worthwhile looking into.

Hype 2: Privileged Access Management (PAM)

As a mainstay of Zero Trust architectures, PAM holds exciting implications in terms of Risk Management but is also put to the test in modern DevOps environments

Privileged Access Management itself is situated on the very right-hand side of the Gartner Hype Cycle, on the plateau of productivity or rather the way to the mainstream. Looking at the PAM maturity curve, you will find a lot of capabilities: from Account discovery – that is, getting the understanding of accounts and providing password management – and Just-in-Time PAM to Service Management integrations and CI/CD automation. I would not say that it makes sense to traverse all these step-by-step. Depending on your IT organization, it might make sense, however, to selectively focus on the capabilities that provide the most value to you.

Privileged Access Management

PAM Maturity Curve

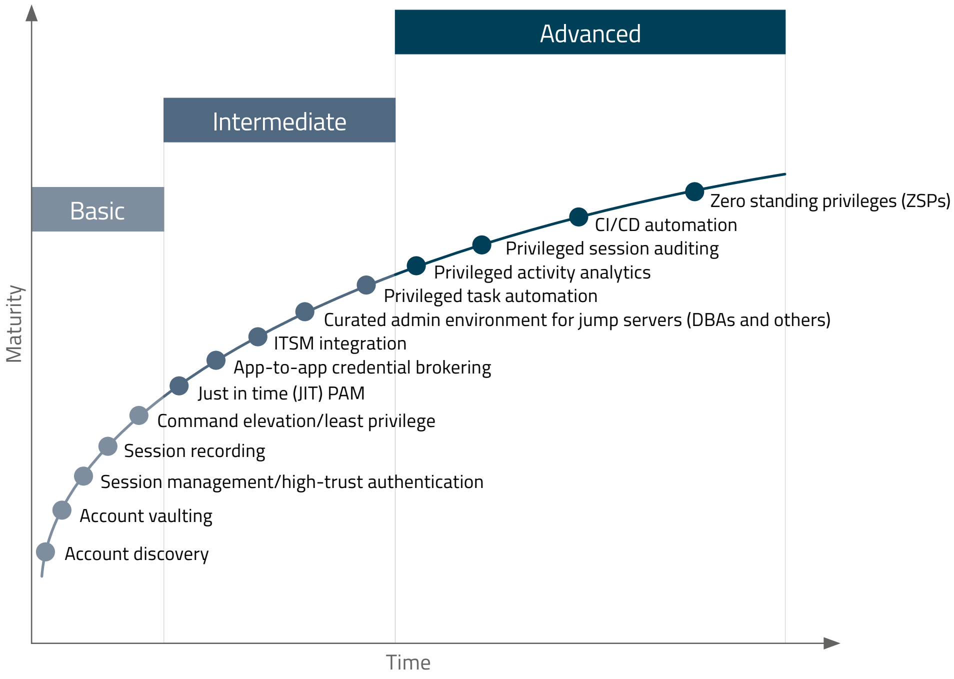

Gartner PAM Maturity Curve

The Role PAM Plays in Establishing a Zero Trust Architecture
Let us now turn our attention to the role Privileged Access Management plays in establishing a Zero Trust architecture. Let us have a look at the approach first and then delve into the specific challenges in the context of PAM projects you should be aware of if you are planning to implement Zero Trust into your IT landscape.

Guiding Principles of Zero Trust

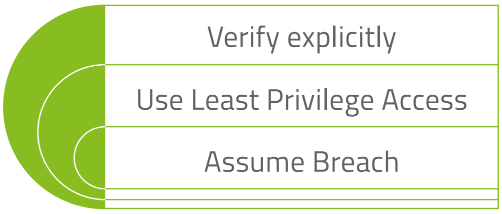

There are three guiding principles when it comes to Zero Trust: First, it is important to never trust and, instead, always verify. Secondly, you should apply the Least Privilege paradigm when granting access and privileges – to the standard users and, even more so, to the administrators, the privileged and highly privileged users. Thirdly, you should always assume a breach, that is, do not get complacent with your security measures and anticipate that your IT infrastructure has either been compromised already or is going to be somewhere down the road and do everything in your power to prepare and mitigate potential damages. Now, how is this approach implemented?

Zero Trust Architecture
High Level Overview

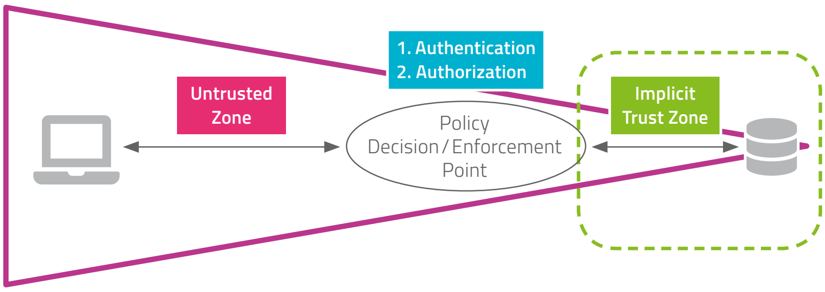

The core idea is to establish a zone of implicit trust, which is as small as possible, for every single resource – everything else is untrusted. In other words: Evaluate and verify requests before exposing a resource. That is the bottom line. Unfortunately, implementation is getting increasingly complex. To stay abreast of these changes, you can leverage a wealth of information within your IT landscape.

Zero-Trust Architecture
High Level Overview

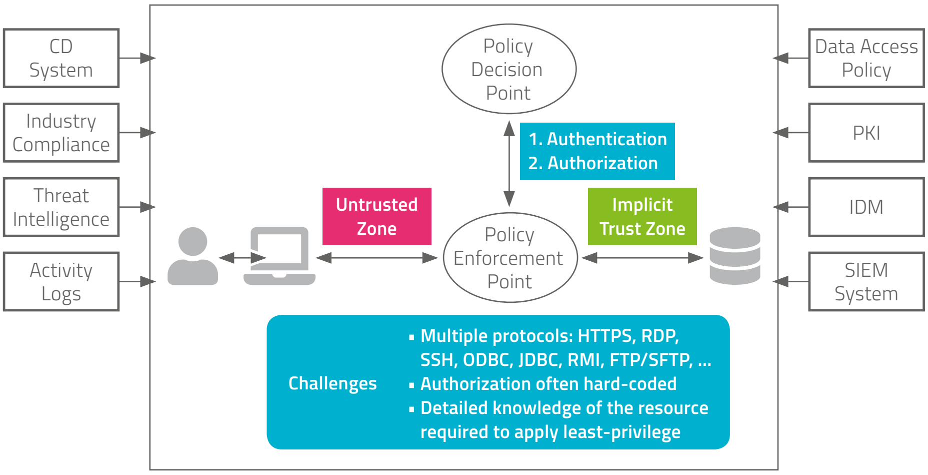

You want to have a layer to enforce access in front of the resources, called Policy Enforcement Point. And you also want a point that is aware of all the policies and makes the decisions to grant or deny access. Of course, authentication and authorization are a crucial major step to provide that access. To do this efficiently, it is not enough to just work with static information based on, for example, the user master data, a role, or access given to a user within an Identity Management system, or a certificate issued by a PKI. You also need to consider dynamic information – that is, not just about the resource and policies applying to that resource but threat intelligence information, activity logs or a device's compliance status information.

The Challenges Privileged Identities Pose

All the above is universally valid, however, not limited to Privileged Access Management. Highly privileged identities pose challenges entirely of their own. For starters, you must deal with multiple protocols. Not just HTTPS but protocols like RDP, SSH and – when it comes to privileged access to databases – ODBC and JDBC. There are tools out there using RMI to make remote calls with proprietary authentication/authorization, making it particularly challenging, and configuration files might be uploaded via SFTP. So, you are dealing with an array of different protocols here. Likely, this will not fit the standard approach for the implementation of a Zero Trust architecture. Another challenge is the fact that, unfortunately, authorization is oftentimes hard coded into systems and there are only a handful of different administrator roles you can choose from. Based on the administrators' daily tasks, in most cases, these privileges are too far-reaching already. With authorization being hard-coded into systems, there's no way to change that, however – which makes for a precarious situation. In case your software vendor provides you with the flexibility to grant fine-grained privileges to administrators – or rather individuals with a higher level of privilege in general – you still need to have in-depth knowledge about the way the system you want to protect operates to apply the Least Privilege paradigm effectively and not restrict your administrators in their day-to-day business.

Identity Risk Management Approaches

The good thing is, there are approaches that help us get a better understanding, mitigate some of these challenges and reduce our attack surface.

Identity Risk Management

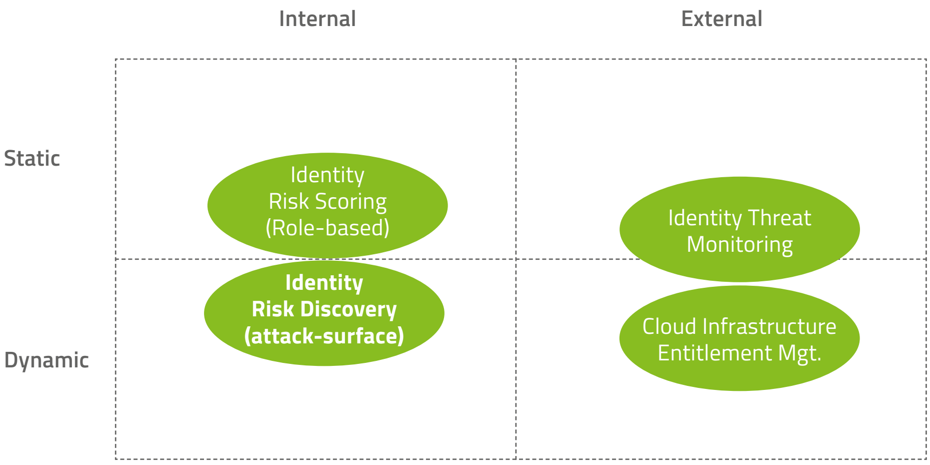
Internal
External
Static
Dynamic
Identity Risk Scoring (Role-based)
Identity Threat Monitoring
Identity Risk Discovery (attack-surface)
Cloud Infrastructure Entitlement Mgt.

While the concept of Identity Risk Management itself is not entirely new, there are some up-and-coming techniques that are worth looking into. One that is receiving renewed attention is Identity Risk Scoring based on privileges. As per the Least Privilege paradigm, you want to avoid accounts with a continually high level of privileges. In that sense, the scoring might not just give you an indication as to the protection required for a specific account but also as to whether you are actually applying the Least Privilege principle. Essentially, it helps you understand not only which accounts but also which roles you should devote your attention to. Focusing more on the external side is Identity Threat Monitoring, that is, getting transparency of the identity landscape based on information of leaked accounts or Darknet information about your organization via dashboard or managed service from your solution provider. And then, there is Cloud Infrastructure Entitlement Management, which focuses on the external side but brings in a lot of dynamic information in the form of live data from Cloud infrastructure providers and resources.

Identity Risk Discovery

A technique I would like to address in a bit more detail is Identity Risk Discovery based on the actual attack surface – that is, not on static information in a PAM system's directory but on live information out there. What does that look like?

Identity Risk Discovery

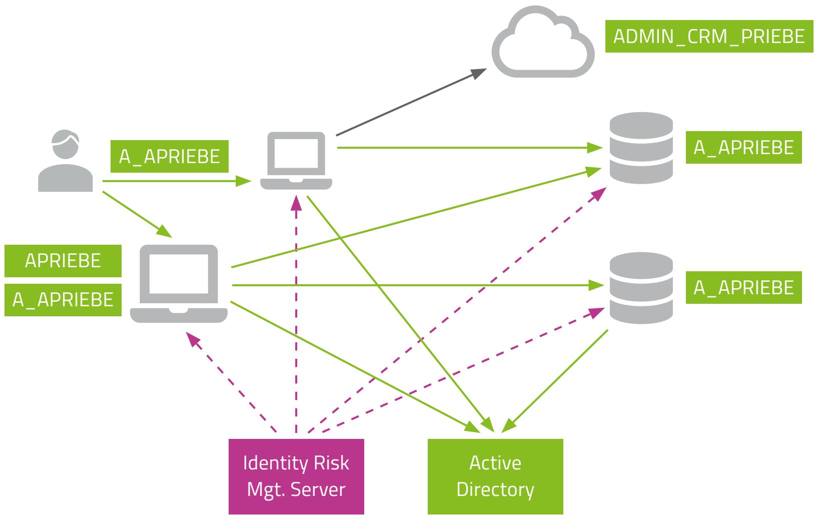

As an administrator, for example, I log into my workstation using my personal account and then log into an administrator account to maintain a given system. Based on Active Directory or the configuration of these resources we understand the privileges. What we typically do not get is a sense of where a specific account is used. Is it just used on one system or on others as well? Perhaps there are other administrator accounts that are not part of the Active Directory, not part of the PAM at all because they are, for instance, Salesforce admin accounts.

The additional value Identity Risk Management services provide is a detailed understanding of where a specific account is used: On which notebook? On which workstation?

Which hashed credentials are there? For example, hashed credentials used for NTLM or Kerberos tickets but also hashed credentials for Cloud access, access tokens, passwords and so on. So, there is a change in perspective – you are not just looking into the repository and the target systems, the resources, but also and especially into the systems on which these credentials are used and stored.

The fascinating thing about this approach is that it mimics the way attackers would infiltrate your organization. Leveraging a security leak, for instance, to compromise one single workstation, one notebook, they would then try to figure out which credentials or rather valuable assets they have at their disposal for their next step and continue to move laterally through your infrastructure until they can escalate an account to domain administrator, for example. Identity Risk Discovery tools follow the same path attackers would, connecting to a given system, taking stock of what is there, collecting data and disconnecting. Subjecting the systems of your choice to this regularly yields you a comprehensive identity risk map, as it were, the kind of map attackers would need to permeate your infrastructure – and based on that knowledge, you can now prevent this. Understanding their usage enables you to make informed decisions about whether widely used accounts should retain high privileges or – for the sake of Least Privilege – if it is more prudent to split them up into multiple lower-privileged accounts, for instance.

Another exciting aspect: As an add-on, many Identity Risk Discovery providers also provision manipulated credentials. As soon as attackers get their hands on and use these, you will be alerted to the fact that a given system is under attack by means of these credentials and can take countermeasures or shut it down – a honey pot approach for identities if you will.

The PAM Journey

Okay, now let us go back to the PAM maturity curve. I picked out three capabilities to exemplify the journey we are taking here.

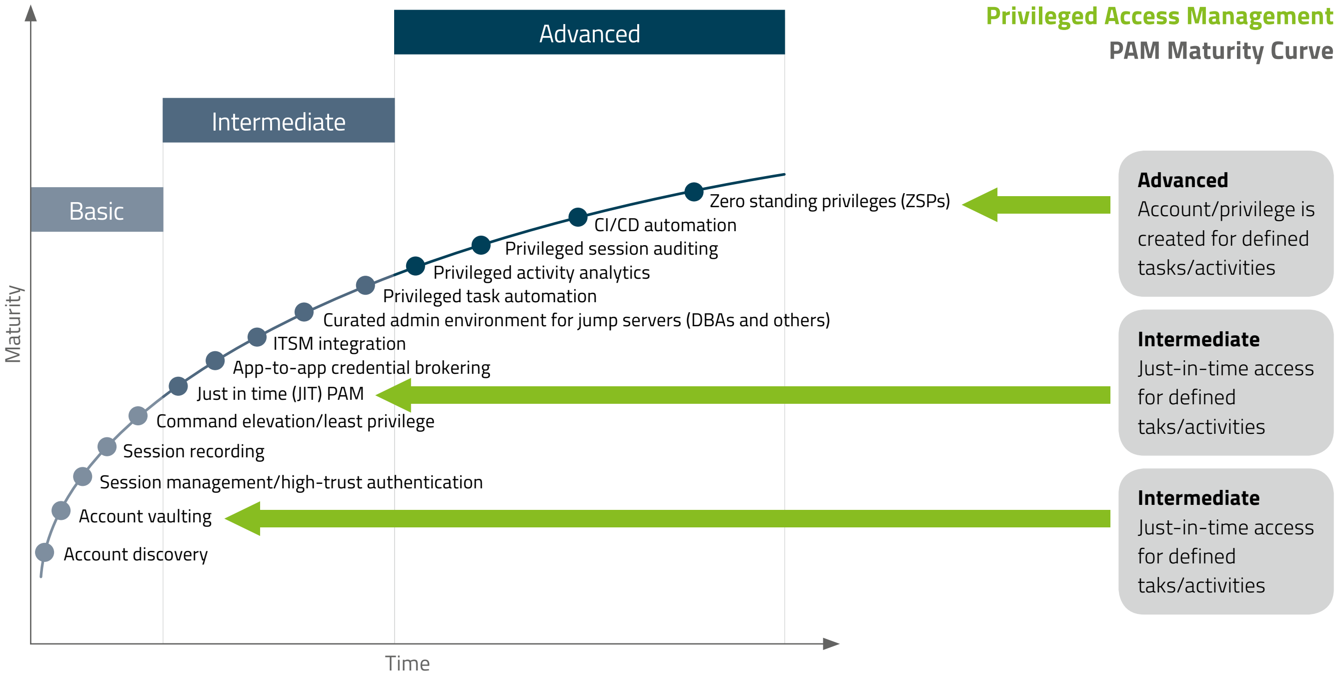
Privileged Access Management

PAM Maturity Curve

Advanced
Intermediate
Basic
Maturity
Time
Zero standing privileges (ZSPs)
CI/CD automation
Privileged session auditing
Privileged activity analytics
Privileged task automation
Curated admin environment for jump servers (DBAs and others)
ITSM integration
App-to-app credential brokering
Just in time (JIT) PAM
Command elevation/least privilege
Session recording
Session management/high-trust authentication
Account vaulting
Account discovery
Advanced
Account/privilege is created for defined tasks/activities
Intermediate
Just-in-time access for defined taks/activities
Intermediate
Just-in-time access for defined taks/activities

At a very early stage, after we have discovered the privileged accounts, we want to vault those credentials, that is create one central place for administrators to retrieve the credentials they need to work with from. The catch: Access is indefinite, unnecessarily exposing a potential attack vector. To address this issue, we transition to Just-in-Time PAM, meaning access is not indefinite but only granted when needed to perform a certain task and then revoked. Even if this account is now no longer accessible, it still exists, though, as do the privileges that come with it – again exposing an unnecessary potential attack vector. This is what the concept of Zero Standing Privileges in an advanced stage of PAM maturity addresses. The core idea is that privileges are provisioned at a given point in time only. So, we are not talking about limited time access – rather, the privileges do not even exist if not required, greatly reducing the attack surface.

PAM from a Risk Management Perspective

If you are wondering whether you really need all these capabilities, whether vaulting is not enough after all, it helps to contemplate this through a risk management lens and get a sense of how the implementation of some of these capabilities' factors into risk mitigation.

Zero Standing Privileges

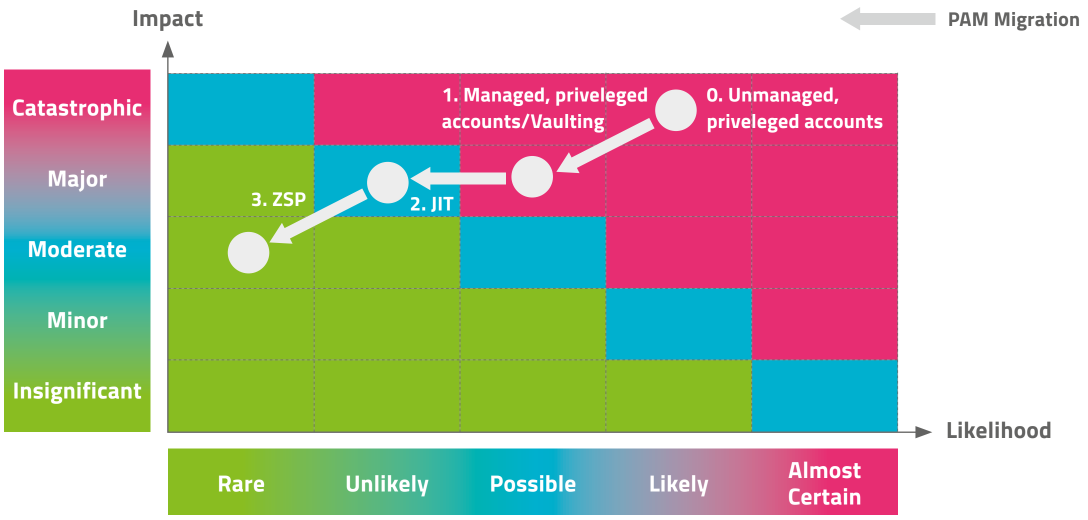

With unmanaged privileged accounts, you run the immediate risk of falling victim to a breach – sooner rather than later, attackers will get their hands on them. Depending on the nature of the compromised account, the impact can be catastrophic. Think of manufactures unable to produce for weeks or even months and the financial impact associated with such a scenario. So, the first step in mitigating this risk would be managing these accounts and having vaulting and access control in place. Transitioning to just-in-time access and, eventually, to Zero Standing Privileges, with privileges being in existence and provisioned at the point in time they are required only, reduces the attack surface as well as the likelihood of a privileged account being compromised even further. At the same time, adhering to the Least Privilege paradigm and having much more fine-grained privileges helps reduce the potential impact of a breach significantly. Combined, this makes for a tolerable risk level.

Multi-Factor Authentication for Privileged Users

Now, let us turn our attention to Multi-Factor Authentication for privileged users and the approaches you can take to implement it even in very complex scenarios.

Security Architecture – Integrated Windows Authentication
Alternative approach

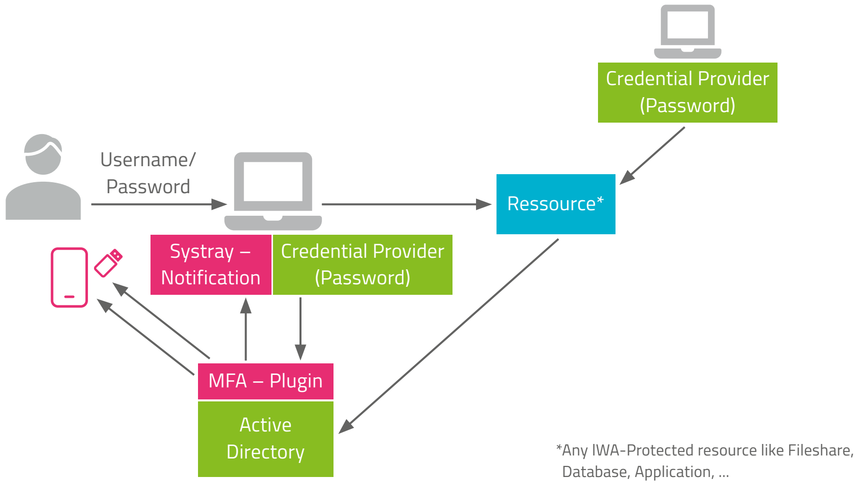

*Any IWA-Protected resource like Fileshare, Database, Application, ...

There are solutions out there that not only focus on protecting the vault but on a scenario in which a ticket-granting ticket is used to get a new session key to access a specific resource. Typically equipping the domain controller with MFA enforcement capabilities, these solutions send out an MFA authentication request to a mobile phone or the system tray of a notebook, prompting users to authenticate themselves. After that, the session key is issued by the domain controller. This helps in scenarios in which you do not feel comfortable providing a password or opening a session for an administrator based on a single factor but really want to have this out-of-band Multi-Factor Authentication before granting access to your organization's crown jewels. Although it makes a lot of sense in this context, this is not restricted to PAM use cases, however. There are a couple of additional benefits you can derive from such an approach. For example, getting an understanding of the service accounts used in your IT landscape and controlling, on a very fine-grained level, for which kind of resources you want to have this strong authentication in place.

PAM in DevOps Environments

Looking at IT landscapes of the past, you would find a plan team, a build team (which has, for good reasons, limited access, limited privileges and, especially, no access to a productive environment) and a run team, which has all the high privileges required to bring IT solutions to life, to operate and maintain them – and Privileged Access Management is typically what facilitates their daily work: deploying newly built IT systems and APIs but also clearing firewalls, implementing the required API gateway policies, onboarding APIs within the Identity Access Management system and so on.

Architecture – Build team and run team

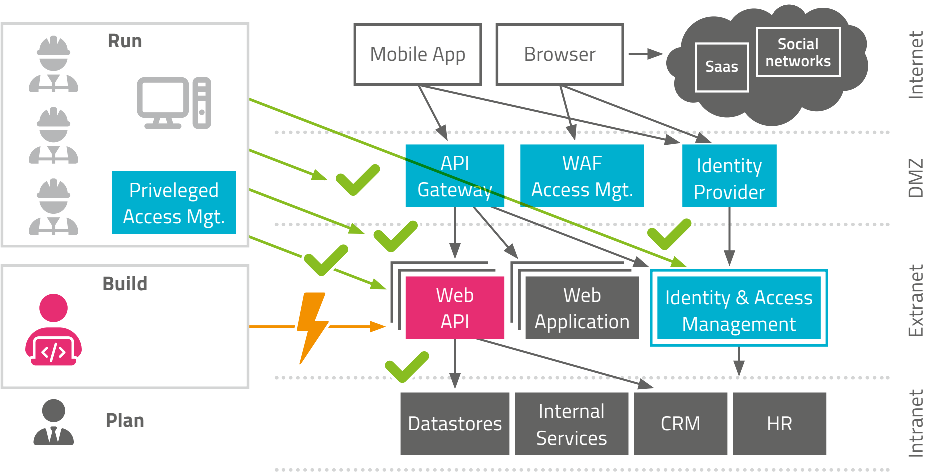

The more contemporary way of going about this, of course, is DevOps – merging build and run into one single team. That does not mean that everybody is doing everything. But the responsibilities are very, very closely related to each other. So, this single team needs the capabilities to perform maybe not all but most of the operations mentioned above.

Architecture – DevOps

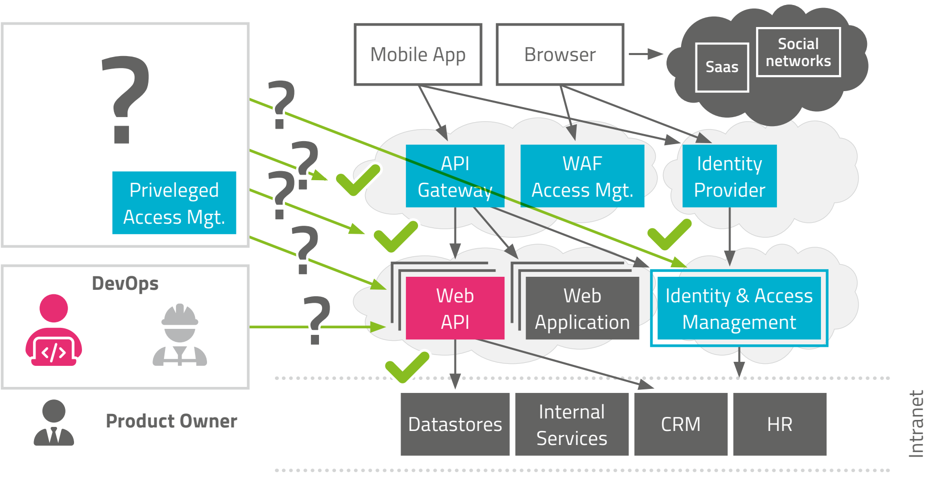

Firewall clearance, you might argue, is not the responsibility of the DevOps team. But maybe there is no firewall or rather no relevant firewall anymore and we are instead talking about a Kubernetes cluster with containers. What is actually exposed is not immediately apparent to a firewall team anymore. Everything is running via HTTPS, through Ingress NGINX and its configurations within that cluster or – if it is a cloud infrastructure – maybe all the rules are in a Terraform script.

Architecture – DevOps & Cloud Challenges

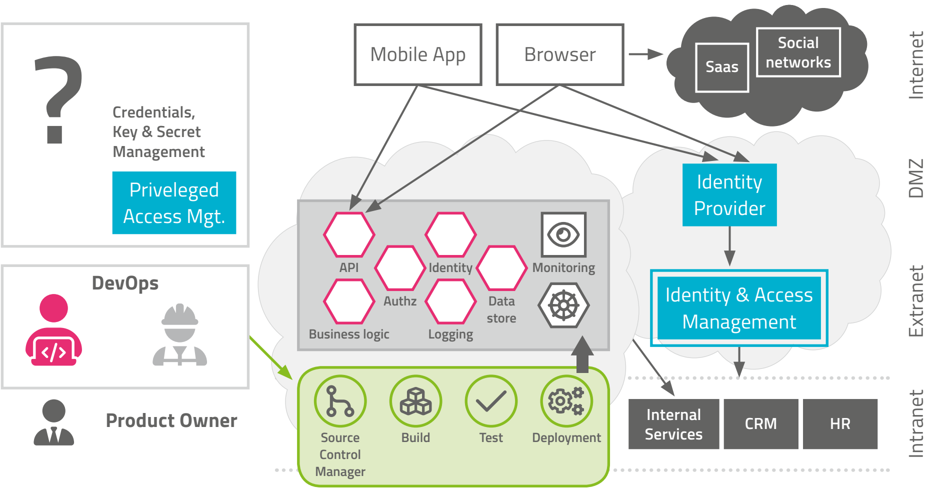

Also, in terms of deployment, we are not talking about SSH logins by administrators but rather completely automated deployment via a CI/CD pipeline. What does that mean from a Privileged Access Management perspective? Unfortunately, it means that things are getting more complex. Who, for instance, decides what kind of service is directly exposed to the internet, why an API is exposed but a data store or logging database is not? At the end of the day, it is code, a configuration file in Git stipulating that an API should be exposed, and a logging system should not be exposed. By implication, that means anybody who has the privilege to modify code in the source control system, in Git, can determine what is exposed to the internet. So, the privileged user is not just the administrator, who might have access to a firewall, but the developer, who is doing exactly the job the firewall administrator did before. In this scenario, you will not readily achieve a level of security equal or comparable to the PBR model without additional provisions.

The next aspect is that a lot of the resources used during the build process are taken from external sources, external repositories like Docker Hub, for instance – and what they are using here is entirely up to the developers. Typically, there is no review anymore, there is no dedicated team screening the package before it is deployed – all of this is happening completely automated. So, additional tooling is required to get a good understanding of what the dependencies are, whether there is a security leak out there, whether the package is on a whitelist and trustworthy or on a blacklist. In other words, we must implement a lot of the provisions from the PBR model into our CI/CD pipeline. The CI/CD pipeline needs to have the credentials to access numerous systems. And these credentials should be managed in a central place to have transparency, be able to revoke them and get an understanding of how often secrets, keys and passwords have been changed, for instance – meaning, in addition to human-to-machine communication, we must take authentication of machine-to-machine communications into account as well, as this is becoming increasingly relevant today. And this can typically be covered by a Privileged Access Management system or by dedicated solutions focusing on managing secrets, keys and so on for our pipelines, for our system-to-system communication.

Micro Segmentation

One way of managing the aforementioned exposure of resources, which is not as straightforward in a containerized world, is adopting a micro segmentation approach (which is also recommended for Zero Trust architectures, by the way):

Micro Segmentation
Alternative approach

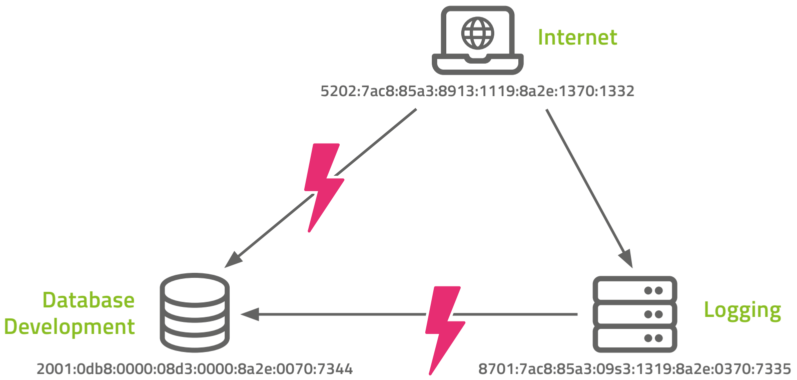
Internet
5202:7ac8:85a3:8913:1119:8a2e:1370:1332
Database
Development
2001:0db8:0000:08d3:0000:8a2e:0070:7344
Logging
8701:7ac8:85a3:09s3:1319:8a2e:0370:7335

Instead of utilizing static firewall rules – which are hard to manage on the one hand, hard to audit on the other hand and would likely not align, for example, with the dynamic IPs in a containerized environment, anyway –, micro segmentation vendors tag and label systems: One system comes from the internet, so it is labelled "Internet". Another system is responsible for collecting log file information, therefore it is labelled "Logging." And yet another system is a database/development hybrid, so it is labelled "Database" and "Development." And now we can formulate policies:

- Network segmentation based on policies and labels
  - Labels:
    "Database" or "Development or "Logging"
  - Policies:
    DENY from "Internet" to "Database"
    DENY ALL from "Logging"
- Agent based approach Instead of perimeters
- Support for multi-cloud-scenarios

Policies stipulating, for instance, "DENY from 'Internet' to 'Database'" – there is really no scenario in which you would want to expose a database to the internet. Or "DENY ALL from 'Logging'" – the logging system should be collecting logs, should be called by the Kibana dashboard or the like but never connect to any other system. By means of this approach, which is working based on agents in the TCP/IP stack rather than physical network segments, we can also support container and cloud scenarios and are not limited to one database, one physical data center.

Conclusion

Within modern Cloud infrastructures and DevOps environments, managing privileged access and identities is becoming increasingly challenging – and organizations need powerful Identity and Access Management solutions to secure and enable their worldwide user bases. Technological innovations like CIEM and Zero Trust 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

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.