ConductorOne

# A Practical Approach To Achieving Zero Standing Privilege (ZSP)

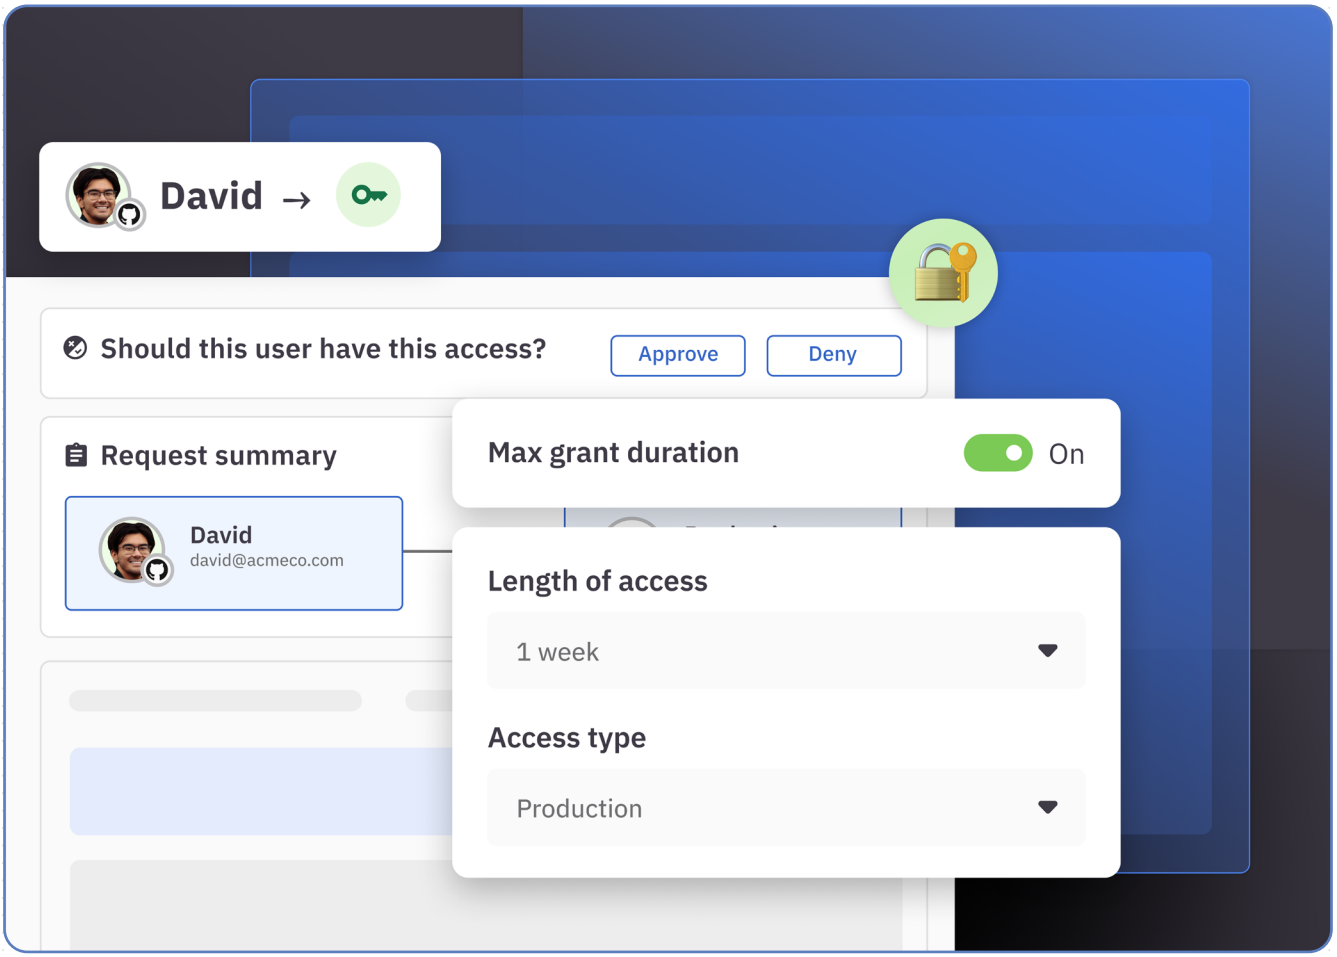

## What is zero standing privilege (ZSP)?

Zero standing privilege (ZSP) is the principle of granting a user or system the minimum necessary levels of access and privilege needed to perform assigned duties for a limited amount of time. ZSP mitigates security risks by greatly reducing the amount of active and privileged access to accounts and systems at any given time, giving hackers fewer potential vulnerabilities to exploit.

ZSP is a core tenant of a Zero Trust approach to cybersecurity, and implementing it has knock-on benefits beyond improved security. These include reducing the scope of access reviews and streamlining compliance efforts. While moving to ZSP is beneficial, it can create logistical challenges for companies to move away from the norm of provisioning birthright and long-lived access.

In this guide, we'll look at the current security landscape to understand why ZSP is an effective solution for protecting hybrid and cloud-first environments and provide a tactical approach to achieving ZSP.

## Why should you make ZSP your goal?

### SaaS and IaaS adoption has exploded

In the past decade, there's been a large-scale and rapid adoption of software as a service (SaaS) and infrastructure as a service (IaaS) applications by organizations of all sizes, across industries. The average company is now using anywhere from ten to hundreds of applications to run their business, ranging from convenient project management and content creation tools like Asana and Figma to business-critical infrastructure like Hubspot, Snowflake, and AWS. IT and Security teams have scrambled to keep up with this explosion of accounts and permissions, but momentum is against them. Depending on a company's size and needs, just a single tool like AWS can comprise hundreds or more accounts and role-based permissions to provision and manage.

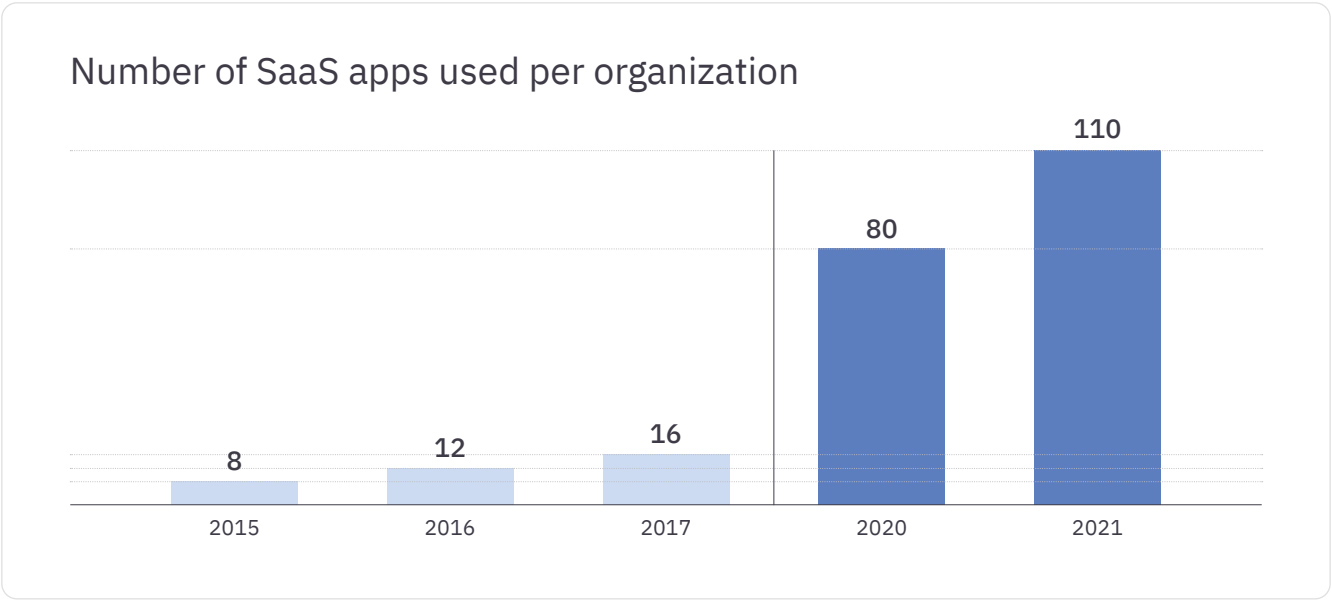


## Efficiency and security are competing forces

Because IT and Security bring their own priorities to access management, they've taken competing approaches to handling this proliferation of applications and accounts. IT teams want to enable business stakeholders with the right technology for productivity, but they're often under-resourced and stretched thin. To keep their workload of access request help desk tickets at a manageable level, they've had to rely on streamlining access management by granting group-based, birthright permissions. If most employees in a given group need access to an application, that application is made available for the entire group, and new employees are given standing access to it from day one. New engineers are provisioned with access to engineering-centric apps and tooling, new marketers are provisioned with marketing team apps and permissions, and so on.

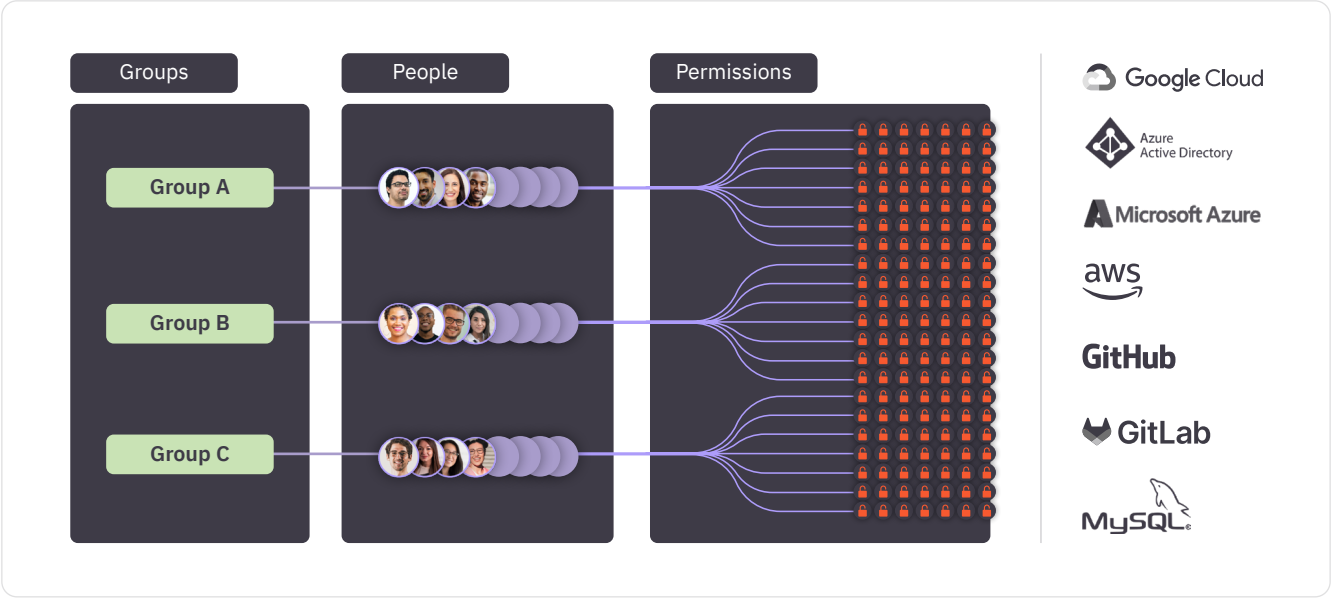

This approach runs counter to the goals of Security teams trying to manage the shifting risks introduced by hybrid and cloud environments. From Security's perspective, birthright access grant models create an expanded "blast radius" of identity attack surface area that is rarely systematically addressed after day one of employee start. The greater the number of users granted access to an application, the larger the risk of identity-based breach from an account compromise. For Security, the safest access is no access at all, which is obviously impractical. And so the next-best pattern is to simply grant access for as long as it is needed.

## Effective permissions and data access are hard to see

Further complicating things, each application uses its own permission model for managing authorization, which makes it hard to contextualize the risks associated with granting any one permission or role. This is especially true when working with large, complex platforms like Snowflake and AWS that employ multilayered hierarchies of roles, each with their own distinct permission ontology, resource authorization enforcement models, and complex inheritance structures.

As permissions sprawl across a range of accounts and systems that include both on-prem and cloud-based infrastructure, it becomes increasingly difficult to get a complete, accurate view of the access users have and how permissions interact, and companies find themselves struggling to catch accounts and overprivileged access that may put their business at risk.

Some account types that pose a security threat:

- **Unused accounts:** Accounts that are not being actively used and are no longer needed.
- **Orphaned accounts:** Accounts that may be provisioned manually and/or are not connected to an identity provider. These accounts often live in systems that are not easily visible and so fly under the radar of security admins.
- **Poorly secured accounts:** Accounts in connected systems that don't meet best-practice security protocols (e.g., a local account that does not use MFA).
- **Non-offboarded accounts:** Accounts that haven't been properly deprovisioned after a user leaves the organization.
- **Overprivileged accounts:** Accounts that have not been properly right sized or deprovisioned after a user changes role or no longer needs access.

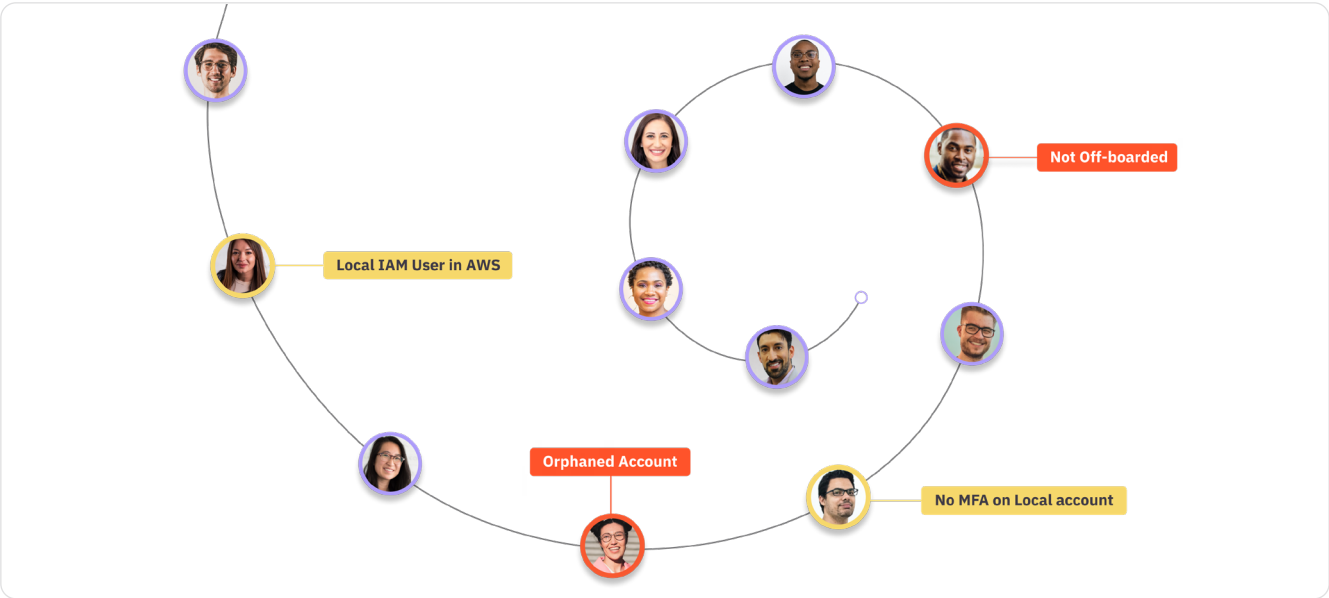

## Hackers aren't breaking in—they're logging in.

*"In a 2022 IDSA survey, 80% of firms reported suffering an identity-related breach in the previous year."*

Over-permissioned accounts naturally increase the risk of identity-based attacks, and 49% of organizations currently have users with more access privileges than are required to do their job. So it's no surprise that 80% of firms surveyed by IDSA in 2022 reported suffering an identity-related breach in the previous year, many of which were the result of "inadequately managed privileges," "compromised privileged identity," or "excessive privileges leading to an insider attack." Hackers are looking for these accounts and know how to take advantage of them. In addition to exploiting directly compromised user identities, hackers can perform brute-force account takeovers and use stolen session cookies from a compromised machine to access downstream systems, among other tactics.

## Traditional PAM doesn't work for the cloud

A traditional approach to privileged access management (PAM) is centered around the use of credentials like passwords to control access to privileged roles and accounts. In a traditional PAM model, a user is given a unique password to gain temporary access to an account. When that user no longer needs access, the system rotates the password so it can't be used again by that or any subsequent user. This model doesn't translate well to cloud-first environments; SaaS and IaaS application privileges are granted based on identity, not credentials—users are assigned roles and permissions in these applications based on their role or group in the organization. This model requires new technical approaches to access management.

# A five-step tactical solution to achieving ZSP

1. **Create a single pane of glass.** You can't protect what you can't see, so a critical first step is to centralize and normalize identity and access data from across your applications and technologies—including cloud and on-prem apps, infrastructure, active directories, and back office apps. Do what you can to pull all application identities, permissions, and authorization data together from your systems under one view, which will allow you to create an overall identity graph of your permissions, roles, groups, and their relationships. While this can be a challenging project, centralizing this information will allow you to put access risks into context and to quickly and efficiently answer questions about identity- and permission-based threats.
2. **Remove unused access.** Looking at usage is a simple way to quickly prune a lot of unnecessary access. Audit accounts for unused or overprovisioned access and remove it. For example, if a user hasn't logged into an account for over 90 days, or if they're not using a particular role within an application, they likely don't need that access on a standing basis.
3. **Operationalize periodic access reviews.** Regular, systematic access reviews allow you to proactively identify accounts that need attention, greatly reducing the instances of orphaned accounts, non-offboarded accounts, and the like. Access reviews are a typical compliance requirement, but they're also just best practice. Create a frequent schedule of reviews and automate them to reduce effort and costs. Even better, perform periodic "micro-reviews" of access based on role or job changes.
4. **Move sensitive access to just-in-time.** As much as possible, sensitive permissions should be provisioned on a need-to-have basis rather than a standing basis. Identify all high-risk access and, as much as possible, move it to just-in-time provisioning with a self-service component. For example, AWS admin access that may have been previously granted as birthright for certain engineers can be moved to time-bound, just-in-time access that has to be re-requested periodically. If the access is not re-requested within the defined period of time, it's automatically removed.
5. **Make it simple to reinstate access.** In addition to removing and limiting access, it's also important to make it easy for users to regain access when needed. Create self-service and automated provisioning that makes granting permissions and privileges back to the user quick and painless. This will allow you to be a lot more aggressive about removing and disabling things in the first place. Users will be more reluctant to let go of things they might need in the future if regaining access is difficult and time consuming.

The benefits of enforcing ZSP

## A systems-wide, consistently enforced ZSP approach to access management hardens your security posture while creating additional benefits across your organization.

With fewer accounts and privileges permissioned overall, the surface attack area shrinks significantly, as does the likelihood of instances of inactive, orphaned, and other exploitable accounts. Though ZSP takes some effort to operationalize, once in place, security teams have the tools to proactively remediate identity-based risks and ultimately avoid costly breaches.

Beyond minimizing security risks, ZSP also results in faster, less costly user access reviews and makes it easier to maintain compliance. Enforcing ZSP also demonstrates to customers and shareholders your commitment to protecting their data. Lastly, well implemented self-service with automated provisioning makes it easier and faster to onboard new team members, facilitate role changes, and offboard departing employees correctly, keeping everyone productive while improving your security posture.

# A Practical Approach To Achieving Zero Standing Privilege (ZSP)

Want to learn more about our identity security platform for modern workforces?

**Get a Demo**

ConductorOne | team@conductorone.com

AICPA SOC