



Intelligent Access

Modernizing the World of Identity with
Just in Time Access



Coauthored by Mario Duarte & Tarun Thakur





FOREWORD

The Principle of Least Privilege

Tarun Thakur • Co-Founder & CEO, Veza

In theory, the principle of least privilege is simple. It is:

“

The principle that a security architecture is designed so that each entity is granted the minimum system resources and authorizations needed to perform its function.

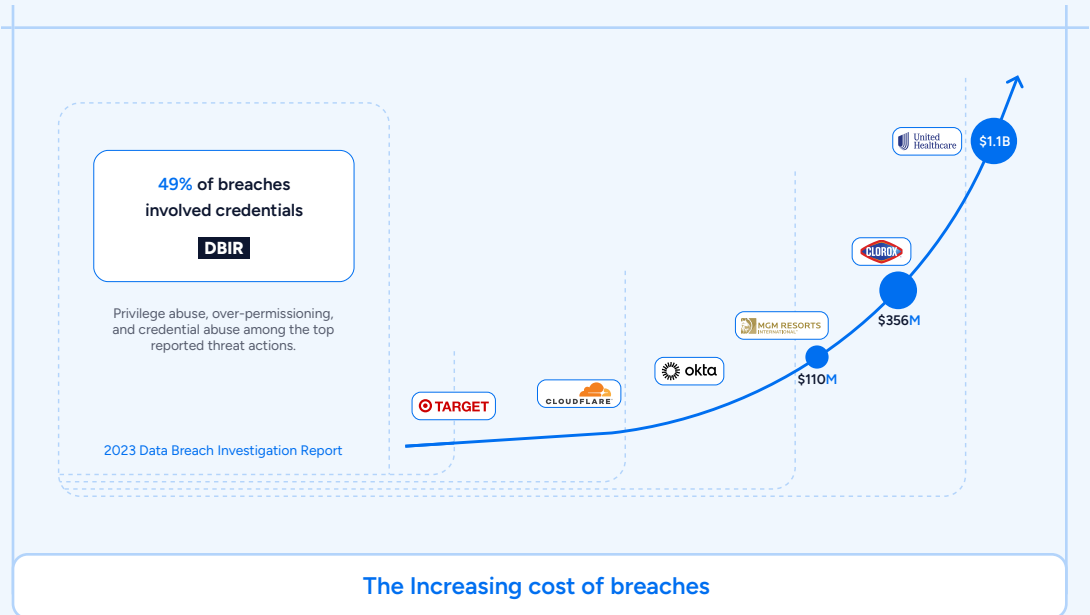
—NIST

Applied to identities (humans, non-humans, machines, etc.), the principle of least privilege means that each identity (including both humans and non-human identities) should only have the permissions it needs to do its work, and no more. Applied effectively, the principle of least privilege promises to protect you from the worst consequences of a compromised identity. For example, if a hacker successfully phishes an employee or a service account (API token, access token, service accounts), the damage they can cause is limited by the associated employee's permissions to enterprise SaaS apps and data systems. The fewer permissions they have, the smaller the “blast radius” from an attack (ransomware, insider threat, credential compromise, etc.).

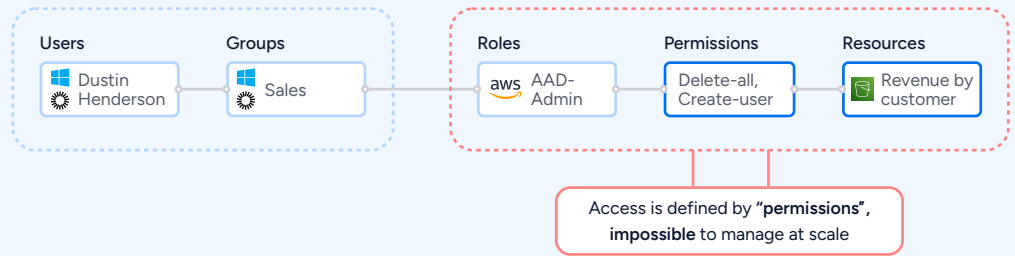
While least privilege sounds simple enough, applying it in the real world is complicated, and getting harder all the time, as the scale and complexity of hybrid and multi-cloud deployments increases. In practice, a “perfect” implementation of least privilege isn't possible. That would mean that no identity ever possessed permissions it didn't strictly need for any amount of time, which isn't realistic for any growing business. It's best to think of least privilege as an ideal to strive for. You'll never get it perfect, but you can probably do better than you're doing now, and any improvement makes you safer from both internal and external threats.

In this short book, we'll dive into the details of why least privilege is so difficult to achieve, and how you can get closer to achieving – and maintaining – that ideal.

Founded in March 2020, Veza was built upon the insight that, in spite of all the existing identity and security tooling available, organizations still struggle to understand “who can take what action on what data”. The [Colonial Pipeline breach in 2020](#) acted as a wake-up call for security professionals around the world. Fast forward to the first [\\$1 billion dollar breach in 2025](#) at Change Healthcare. These data breaches — including Target, Okta, MGM, Microsoft, and more – have brought cybersecurity’s greatest challenge to the forefront: that identity is the weakest link in modern security posture. The current identity access infrastructure is exceedingly complex and difficult to secure, leading enterprise security teams to struggle to respond and keep up with the pace and sophistication of attacks.

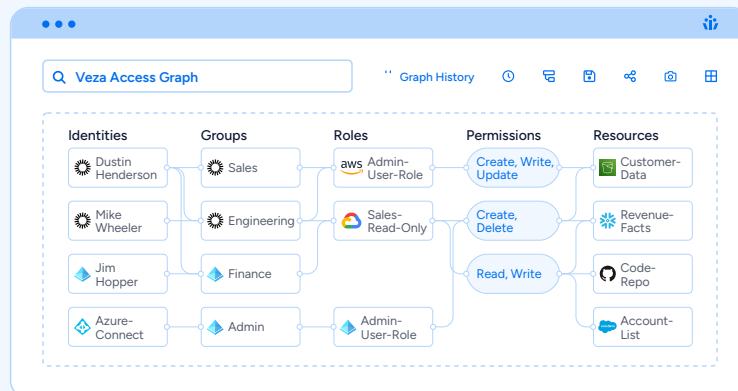


When founding Veza in 2020, my co-founders and I had the insight that all the major identity tools (e.g., Microsoft Active Directory, Okta, Google Workspace, etc.) are just directory services with users and groups. Directory services give limited-to-no insight into access permissions, the purest form of truth of identity. Without the ability to leverage access permissions, these directory services act as mere abstract “directory stores” with which security teams can make semi-educated guesses, largely based on static names and user/group descriptions. It was clear to us that the time had come for a massive transformation of the existing identity tools – identity needs to move from IT to becoming a critical foundation of the Security ecosystem, where the future of identity platform needs to be designed on entitlements and permissions, not users and groups.



All existing identity tools are designed on "Directory Services"

Access permissions universally define who can take what action on what data, but the permissions themselves are system-specific and have no common language. We realized we must organize all these permissions in a canonical data model that associates identities to their "effective" access: normalized to the common actions of create, read, update, and delete. Only then could we truly start to work toward building an identity platform that addresses a wide breadth of use cases, and enables companies to achieve the principle of least privilege.



- ✓ All Systems (Cloud, On-Prem)
- ✓ All Identities (Humans, Machines)
- ✓ True Permissions
- ✓ Standardized
- ✓ Automated

Veza Access Graph

After four years of diligently building toward our vision, we believe that Veza's Access Platform will define and lead the modern Identity Security category by empowering organizations to definitively answer the question, "Who can take what action on what data?" We aim to deliver next-generation identity products for critical business needs:

IGA (Identity Governance and Administration)

Cloud PAM (Privileged Access Management)

ISPM (Identity Security Posture Management)

NHI (Non-Human Identity Management and Security)

SaaS Access Security

Data System Access

Together, these products will enable businesses to achieve least privilege across all their systems. One of Veza's key differentiators as a platform is the breadth of coverage and integrations across all different classes of target systems. Veza helps achieve least privilege for cloud infrastructure platforms (like AWS, Google Cloud, Azure, and OCI), identity systems (like Okta), structured data systems (like SQL Server), unstructured data systems (like SharePoint Online), SaaS apps (like Salesforce and Github), on-premise systems, and custom applications.

Snowflake is a great example of how Veza works on an individual system. Snowflake powers many modern enterprises, but traditional identity tools have never been able to effectively wrangle and secure access to data within.



Find and fix over-privilege for **Security Engineering** teams

Goals

Gain actionable insights and remediation recommendations around violations of policy, security gaps, or dormant permissions for humans and non-humans

Products

Access Search & Intelligence

NHI Search & Intelligence

Activity Monitoring

Budgeted Projects

ISPM / ITDR / CIEM

NHI Security

Data Systems Access

Cloud PAM

SaaS Security



Review and provision/de-provision access through modern IGA tools for **Identity** teams

Goals

Provide complete & modern IGA product suit, OR compliment legacy IGA deployment

Products

Access Reviews

Lifecycle Management

Access Requests

Budgeted Projects

Augment / Compliment IGA

Greenfield IGA

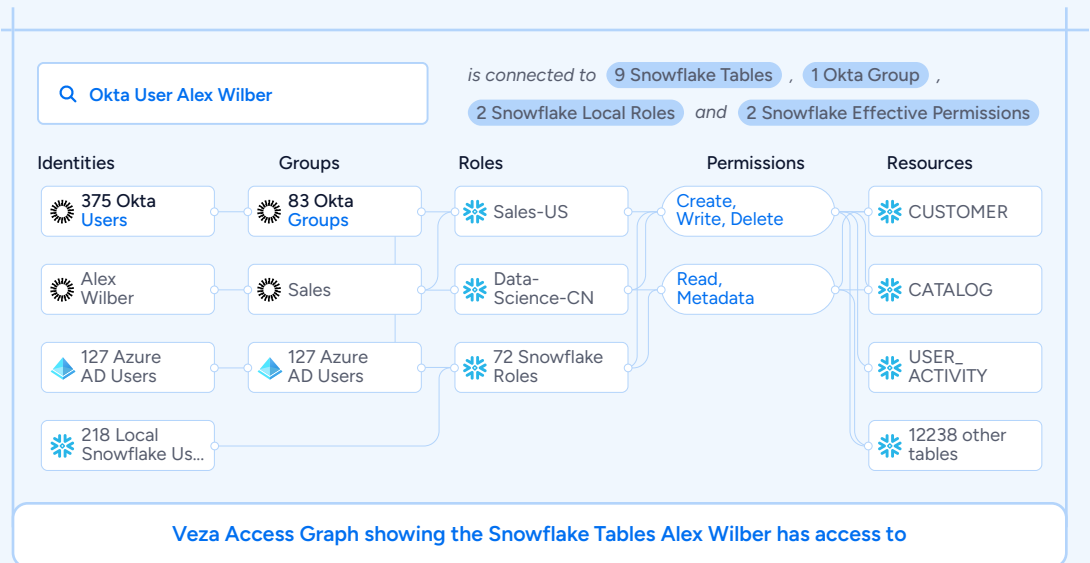
Replace IGA

The Identity Security Use Cases

Can You Tell Who has Access to What in Snowflake?

Over the past decade, Snowflake has grown to become the default cloud-native modern data solution for storing and querying enterprise data. Snowflake's thousands of customers run more than five billion queries every single day. If you're among those ten thousand, Snowflake is probably your single largest repository of sensitive data, potentially including telemetry about customer behavior, PII, and payment transactions. As more and more services build on top of the Snowflake data warehouse, managing access to that data only gets harder to scale.

With over half of data breaches involving identity, the most important action you can take to secure your Snowflake data is to establish tight access control and apply the principle of least privilege to users and roles in Snowflake. However, most organizations struggle to achieve this. They have no idea who really has access to what data in Snowflake or whether that access is being used. Let's look at why this is, and how Veza can help restore visibility to permissions in Snowflake.



Why Don't You Have Visibility into Permissions in Snowflake Today?

Organizations attempting to adhere to the principle of least privilege and follow identity security best practices in Snowflake are confronted by a fundamental lack of visibility into access at the object level, such as to databases, columns, tables, views, or schemas. In other words, they don't really know who can perform what action on what data in Snowflake. And if you can't see who has what permissions, you can't hope to meaningfully apply the principle of least privilege. This lack of visibility into access in Snowflake is driven by four key challenges.

01 Access Controls (RBAC) Complexity

In some ways, Snowflake's system of role-based access control (RBAC) may seem straightforward: each user can be allocated one or more roles, each of which confers a set of permissions to act on objects in Snowflake. However, actually applying RBAC to manage access to data in Snowflake at scale becomes complicated in practice for several reasons:

Snowflake supports more than fifty distinct types of privilege and over a dozen object types. Understanding even simple privilege statements in Snowflake requires specific and extensive expertise in Snowflake's access model.

Users can be assigned more than one role, and any two roles can overlap significantly in the privileges they grant.

Roles are hierarchical: that is, role A can be granted to B, and role B can be granted to role C. So any user assigned role C will also have all the privileges of roles A and B. Role hierarchies are flexible but also make it easy to assign unintended extensive privileges quickly, without the person assigning a role being aware of how much access they're actually granting.

Accidentally nesting a highly-privileged role, like SYSADMIN, under a role with a safe-sounding name like MARKETING can spread dangerous privileges to many identities. This is a common case where "ReadOnly" and non-sensitive roles acquire accidental write or sensitive read permissions due to nesting that go undiscovered. These kinds of mistakes are difficult to discover and often remain undetected until they permit access to an attacker.

02 Scale

Both the amount of data and the breadth of use cases supported by the Data Cloud means that security and governance teams are managing many more identities with access to many more resources than was normal in the on-prem era. An enterprise Snowflake deployment might include thousands of users with access to hundreds of thousands of tables, schemas, and views. The old paradigms for securing access, like the manual quarterly access review, simply can't keep up with the scale of the modern data warehouse. This is only becoming more true as non-human identities proliferate to outnumber their human counterparts on average by **17:1**.



Non-human identities outnumber human identities on an average of 17:1

03 Siloed identity and access data

Most organizations manage access to the Data Cloud via an Identity Provider (IdP) like Okta or Azure AD (Entra ID). For example, it's common for multiple users in the IdP to share access to a single local User in Snowflake. This makes it hard to determine who really has access to what in Snowflake. Teams granting access to Snowflake in the IdP don't know what access they are really giving users at the object level. Meanwhile, admins in Snowflake, looking at the permissions of a local role or user, have difficulty determining how many federated identities get those permissions or would be affected by a change.

04 Data teams vs. Security teams

One of the most significant challenges in maintaining least privilege in Snowflake isn't directly related to Snowflake's access control system at all. It's simply that, in a fast-paced working environment, IAM teams are under pressure to enable data teams that depend on access to the Data Cloud for their work. This focus on enablement has two important consequences:

- The need to grant access immediately will always be more urgent than the need to remove access that is no longer needed, leading to a tendency for identities to accumulate privileges over time.

- Tight timeframes for granting access may prevent IAM teams from taking the necessary time to find the most privilege-efficient way to meet an access request or even from fully understanding what the outcome of granting a new role will be at the object level.

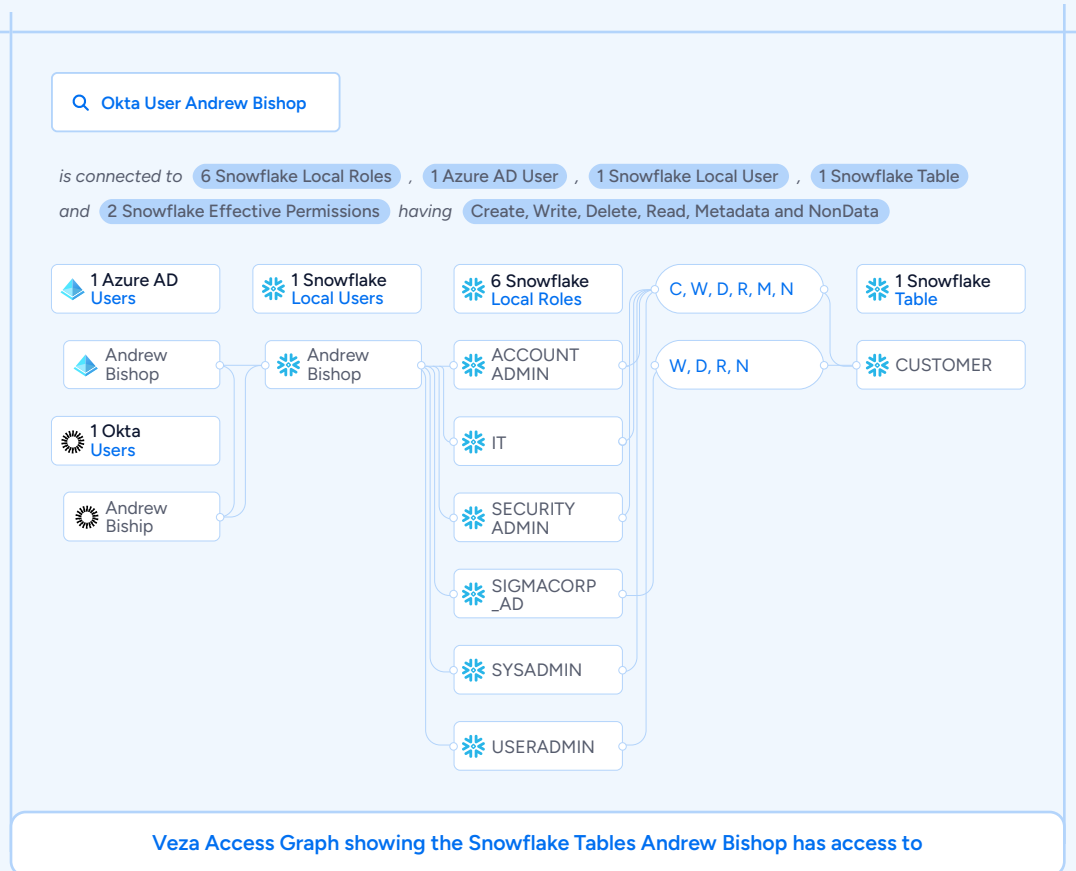
The Cost of Not Knowing Who Has Access to What in Snowflake

These challenges make maintaining best practices and least privilege in Snowflake difficult. This can lead to a number of significant risks, including credential compromise, ransomware, data breaches, compliance failures, technical debt, excess spend, and more.

How Veza Can Help

Access Visibility: Who has access to what in Snowflake?

Veza is powered by its Access Graph, which gives organizations the ability to visualize access relationships between all identities and systems by connecting users, groups, roles, and permissions. This means that you can truly know who has access to what in Snowflake. For example, below, you can see Veza's visualization of all identities with access to the CUSTOMER table in Snowflake, including both local users and federated identities in Okta. Veza shows not just who has access but what kind of access – create, read, update, and delete – and how they have it, including showing the effect of role hierarchies.



Access Reviews

Veza’s Access Graph lets you know the full permissions of any identity on all data objects in Snowflake. This means that Veza can make your compliance simpler, cheaper, and more robust. Veza’s Access Reviews product can automatically compile, schedule, and assign access reviews based on users’ granular permissions. Decision makers have access to all the context shown by the access graph to help them make the right decision. Meanwhile, integration with your ITSM tools like Jira and ServiceNow allows you to automate follow-ups to make sure that rejected access is actually removed.

Snowflake Local Users UAR Q2

Completed Rows: 0/5

0 Approved

0 Rejected

5 Need Review

Source	Permissions	Destination	
Name	Permissions	Name	Type
Snowflake SigmaCorp	MetadataCreate, MetadataRead	ANDREW BISHOP	Snowflake User
Snowflake SigmaCorp	MetadataCreate, MetadataRead	JOHNBAON	SnowflakeUser
Snowflake SigmaCorp	MetadataCreate, MetadataRead	KENNETH	SnowflakeUser
Snowflake SigmaCorp	MetadataDelete, MetadataCreate, ...	SNOWFLAKEADMIN	SnowflakeUser
Snowflake SigmaCorp	MetadataCreate	THOMASSTEWART	SnowflakeUser

Approve

Reject

Sign Off

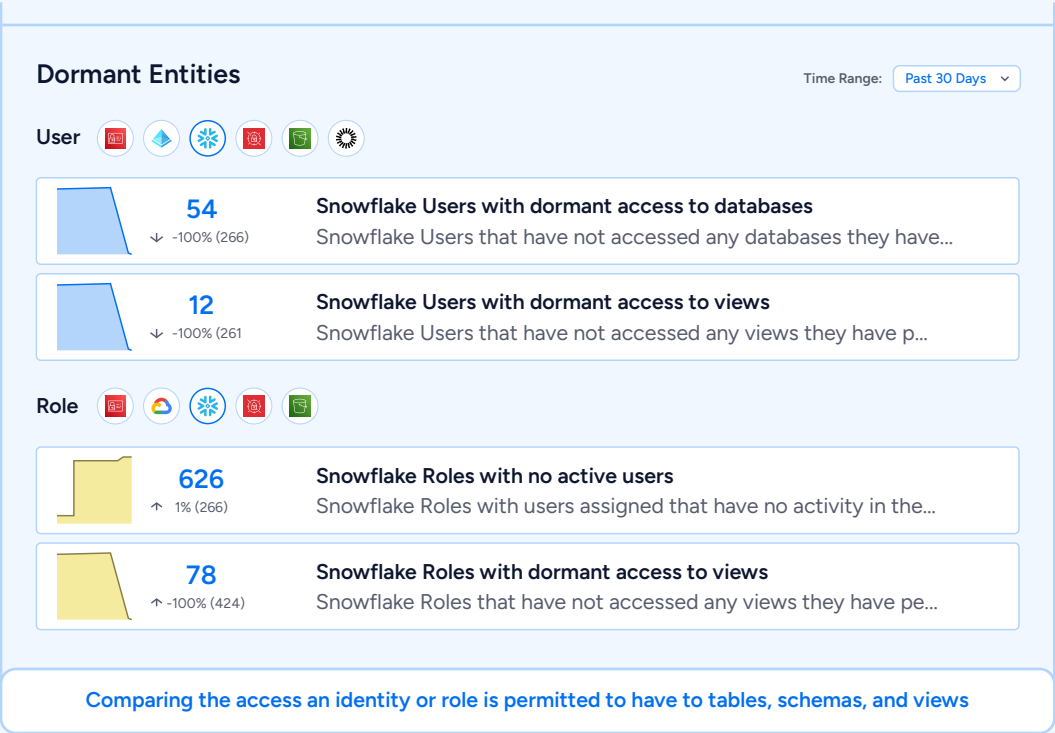
Add Note

Automatically compile, schedule, and assign access reviews based on users’ granular permissions

Activity Monitoring for Snowflake

The principle of least privilege states that any identity should have only the privileges it needs to do its work. For identity teams, this poses the challenge of knowing which privileges are needed and which are not. A good start would be to know if an identity is actually using the privileges it has. With Activity Monitoring, Veza compares the access an identity or role is permitted to have to tables, schemas, and views in Snowflake to the access it actually uses over a period of time. This comparison is expressed in the form of an Over-Provisioned Access Score (OPAS). For example, an identity with access to 10 tables that has only actually used three tables in the last 30 days would have an OPAS of 70%.

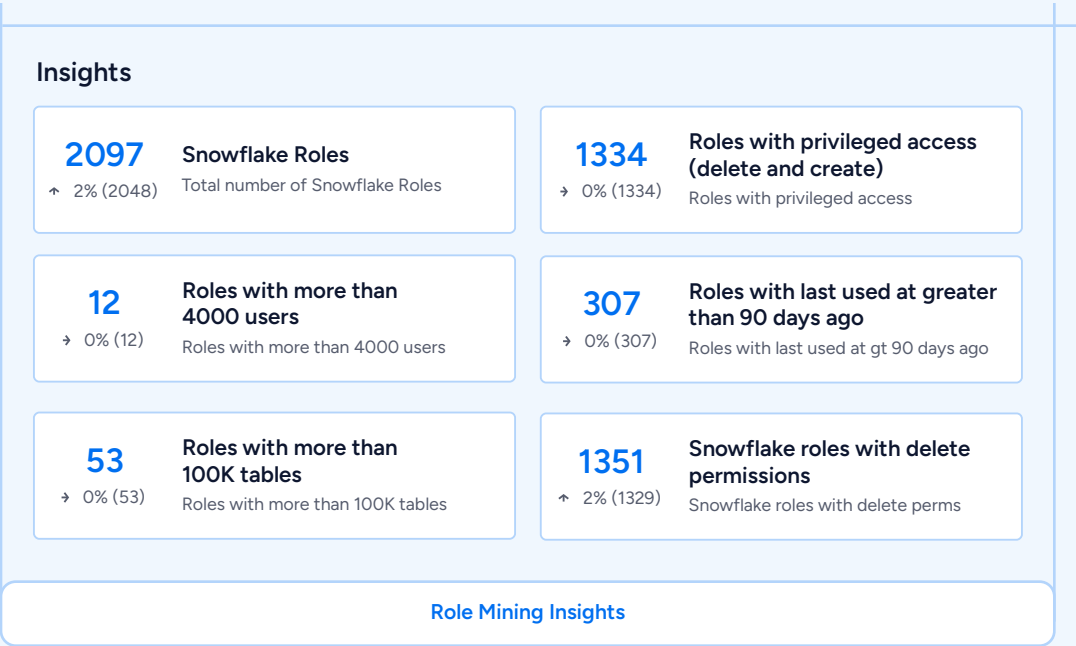
In practice, Veza has found that almost all Snowflake users have an OPAS of higher than 80%. This means that by identifying unused access, Activity Monitoring allows you to reduce the total amount of privileges in Snowflake by up to 80%, making it one of the single biggest actions you can take to make least privilege a reality in your organization.



Veza Role Insights & Role Mining for Snowflake

The visibility into granular permissions provided by Veza’s Access Graph empowers you to restore RBAC best practices and role hygiene in your Snowflake Data Cloud. With a wide range of out-of-the-box insights and the ability to create your own custom queries, Veza’s role mining capabilities allow you to:

- Remove roles that are dormant, duplicative, or used by only a few identities.
- Split up overly generic roles with too many users into more targeted and less privileged roles.
- Identify and untangle complex role hierarchies with multiple layers of inheritance.
- Easily right-size over-permissioned roles by referencing activity monitoring data.
- Triage roles for review by identifying roles with a high blast radius (access to a high proportion of your total schema) or powerful privileges like write and delete.



Snowflake Role Recommendations

Once you've cleaned up your RBAC and weeded out dormant, duplicative, and excessively permissioned roles, Veza's role recommendation tool gives you an easy way to apply the principle of least privilege to respond to access requests going forward.

Let's say Adam requests read access to the CUSTOMER table in Snowflake. Veza analyzes all available Snowflake roles and recommends the role that grants Adam the least privilege access.

Role Recommendations

Integration

Snowflake

User

ANN.TIMES@EXCORP.COM

Resource Type

Snowflake Table

Resource Name

CATALOG_SALES

Effective Permissions

DATA_READ

Roles Filter Query

Snowflake Roles

Recommendation 1

DATA READERS

Least Privileged

Resource Type	Current Access	New Access
Snowflake Table	0	2

Change

Effective Permissions

+2

Create, Delete Metadata

System Permissions

All, Alter, Alter Table, Add Column...

Recommendation 2

ACCOUNTADMIN

Resource Type	Current Access	New Access
Snowflake Table	0	23

Change

Effective Permissions

+23

Create, Delete Metadata

System Permissions

All, Alter, Alter Table, Add Column...

Apply the principle of least privilege easily with Role Recommendations

The Veza Advantage

Securing your sensitive data in Snowflake means establishing meaningful access controls, and you can't do that if you are blind to permissions – to who can take what action on what data. With the power of Veza's Access Graph behind you, you can:

- Remove the risk created by excess privilege and misconfigured identities.
- Ace your compliance obligations while spending less time and money on manual reviews.
- Empower your teams with the access they need when they need it.



Just-In-Time Access

Mario Duarte • Former VP of Security, Snowflake

SaaS and cloud applications have revolutionized the way we do business. Over the last decade, the proliferation of SaaS and cloud applications has come to feel at times like a blessing: organizations can now run more efficiently than ever while enabling a remote workforce. However, security teams are also cursed – they now have to secure access across the enterprise for tens of thousands of users, human and non-human alike. Most wouldn't trade the significant, positive impacts on our working models, like increased productivity, efficiency, and speed, but it cannot be denied that the recent migration to the cloud has created the "Wild West" of access.

Securing access across the enterprise has never been more challenging. Every day, companies issue credentials that increase their potential identity attack surface, and the scale and complexity of entitlements have exploded in the last decade.

The recent introduction of artificial intelligence (AI) has made waves in the tech landscape as well. Similar to the opportunities created by SaaS and the cloud, there is enormous potential to change the way we do business using AI – and we are also hard-pressed to protect against the associated attack vectors.

With the rapid advancement and adoption of AI, concerns around identity and access have been growing. Although AI is the topic everybody is talking about now, it isn't a new issue. Attacks by groups like the Lapsus\$ group have magnified how AI can further the agenda and success of social engineering attacks.

The rapidly changing landscape created by the rise of SaaS, cloud, and AI has served as a catalyst, forcing us to improve our identity and security.

Just-in-time Access for Business



Business SaaS Applications

Business SaaS applications have shifted how attacks happen. Data is stored directly in these applications, with thousands of users and systems being allowed access. Rather than trying to compromise the integrity of the core security systems themselves, bad actors can now focus their attacks on hundreds or thousands of individual employees, who already have the keys to enter these systems. Accessing credentials eliminates the need to break into the cloud infrastructure itself and increases the overall attack surface.



The Increasing Attack Surface

The increased attack surface business SaaS applications offer leads to an exponential access problem. Think of the systems your organization has available to access, and then multiply that by the number of employees. Then, consider the non-human identities in the mix – systems that need access to other systems – to read, write, and execute. This is a massive attack surface.

Start by imagining your core databases and data warehouses in the cloud. You could imagine having only 5 or 6 roles for the whole company – not too difficult. The issue comes with customization – it's difficult to imagine ALL the possible needs of someone in the business. Someone is assigned a role based on their team, like marketing or sales, but their role's access may not allow them everything they need to do their job. They are then given special, custom access privileges beyond their assigned role. There's a demand on administrators to make these types of changes happen quickly and eliminate blockers for individuals to complete their duties. The individual may then be granted more access than they need across multiple systems because it's a complex, time-sensitive issue.

Now, imagine how often this process is repeated over time across your entire company. This is an exponential problem, inflated by the traditionally permanent assignment of permissions—in practical terms, we almost never downgrade access. Our attack surface only grows, with an increased amount of credentials being available with difficult-to-track permissions.

Similarly, SaaS apps have streamlined how users can get additional permissions. Ease of use means that the admins may no longer need to have a technical background, let alone a security background. Individual teams have ownership over their apps, serving as admins, and are able to add users quickly and easily without fully understanding what those users then have access to. Permissions to these apps are then infrequently reviewed, with fewer users being removed than added each cycle — and the attack surface grows again.

The concept of Just-in-Time (JIT) access – allowing limited access in a restricted time frame – would alleviate some of the risks created by these industry practices. Just-in-Time (JIT) access is not yet the industry standard, however, due primarily to a historical gap in our ability to accurately track, at scale, what individuals and entities have what access to what data. It has been difficult to look at things holistically and ascertain how much access an individual has and what degree of that access is actually needed to perform their role. There have been advancements in this space that businesses must take advantage of to secure their systems and lessen the attack surface.



Access is "Access"

There is no difference between an app that has been granted access to your permissions and an attacker that has stolen your credentials — access is access.

Admins are being asked to move faster than ever to compete in business. They are being forced to figure out access across potentially hundreds of apps and thousands of people. It's an issue of scale, and over-provisioning is common, not out of carelessness but because of an inability to accurately communicate who needs access to what data.

Enduring access to one system that has access to other systems – even read-only access – creates the same issues and complexities in securing access for individuals. Just-in-Time access must be applied to both human and non-human identities, if least privilege and a decreased attack vector is to be achieved.



Second Bold Act in Identity

There has been slight progress when it comes to Just-in-Time access, primarily in cloud infrastructure or highly definable environments among highly technical people. For example, a cloud engineer might need to modify the cloud environment. In order to do this, he or she needs some root-level admin access to the cloud environment – but only needs it once: the perfect use for Just-in-Time access.

Just-in-Time access needs to be scaled for broad and generalized use. The greatest advancements we've seen for Just-in-Time Access have come from Veza, and their approach to next-generation IGA – the bold second act Identity needs.

By going beyond users and groups to understand effective permissions, Veza helps companies find and fix risky permissions and policy violations. It secures access to data in virtually any system in the modern enterprise stack, whether on-premise or cloud. And it does so for all identities, whether human or non-human.

Veza provides the ability to visualize who has access to what, monitor privilege drift, and investigate identity threats with access search and access intelligence. Access reviews can be automated to collaborate with the business on smart access decisions. You can also manage provisioning and deprovisioning throughout the identity lifecycle.

Veza supercharges the security for Just-in-Time applications, bringing the tools needed to manage the access evolution generated by SaaS and cloud apps. Businesses need Just-in-Time access to achieve least privilege and step into the future of identity.

Meet the Authors

A serial entrepreneur, Tarun Thakur is a Co-Founder & CEO of Veza, the identity security company. Tarun co-founded Veza in 2020 to address one of the toughest issues in enterprise security today: who can take what action on what data. Veza has established a robust customer base consisting of Fortune 500 and global enterprises like Intuit, Blackstone, Wynn Resorts, Sallie Mae, KKR, Expedia, CrowdStrike, and many more. Veza has raised \$120M+ in venture financing from leading firms such as Accel, Google Ventures, Norwest, True Ventures, and Ballistic Ventures.

Prior to Veza, Tarun was Founder/CEO of Datos IO, and has extensive B2B SaaS experience at companies including IBM Research, Symantec, Data Domain (acq. EMC), and others. Tarun graduated from Duke University with his MBA in Strategy and Marketing. Tarun holds 20+ patents in various fields of distributed systems, security, and next-gen storage and compute architecture.



Tarun Thakur
Co-Founder & CEO,
Veza

Mario has 20+ years of experience as a security professional working in the tech, retail, health care, and financial sectors. He has built and managed security teams and developed and implemented security programs for private and public organizations. He serves as an advisory board member at several cybersecurity companies as well as an investor for early stage startups in the cybersecurity space.



Mario Duarte
Former VP of Security,
Snowflake



The Identity Security Company