Teleport

Defending against Identity Provider (IdP) Compromise

Implementing Infrastructure Defense in Depth - keeping critical systems and data secure in the event of identity provider compromise

White Paper

Contents

Teleport partnered with security researcher Doyensec to explore the current state of the third-party identity management security market and steps that companies can take, using Teleport, to implement an additional layer of security in order to reduce the impact of a compromised SSO provider. The paper summarizes Doyensec's findings, including a Hardening Checklist for defense against IdP compromise, as well as commentary from Jack Poller, Principal Analyst of Paradigm Technica. Customers may request Doyensec's full research report from their Teleport account manager.

Introduction

As organizations increasingly adopt Single Sign-On (SSO) as the central point of authentication, Identity Providers (IdPs) become a focused point of attack for threat actors to gain access to critical infrastructure and sensitive data. Recent incidents have highlighted a growing threat landscape targeting these identity ecosystems, posing significant risks to the security posture of both businesses and individuals.

> Organizations must take action now to harden their infrastructure against compromised identity provider scenarios.

This paper introduces Infrastructure Defense in Depth, an approach to infrastructure security that companies can adopt using Teleport to keep their critical systems and data protected even in the event that their identity provider is compromised, using the capabilities and configurations outlined in this paper.

The Threat Landscape

Recently, there has been an unprecedented rise in threat actors targeting Identity Providers (IdPs). As organizations increasingly rely on centralized identities for user authentication and access control on a wide range of business-critical services, the compromise of an identity provider poses increasingly severe risks to downstream infrastructure.

The Single Sign-On (SSO) provider is chosen as an entry point by attackers specifically because of its trusted highly privileged position managing authentication within the target business. Common patterns employed by attackers include: social engineering, broad-based or spear phishing campaigns, bribing employees for 2FA codes, prompt-bombing, credential stuffing, session hijacking, password spraying, access tokens leakage and counting.

The blast radius of potential attacks heavily depends on the number of Service Providers (SPs) that trust the identity provider to authenticate and authorize internal actions, as well as the level of compromise itself:

- **IdP Vendor Compromise:** This scenario delineates all cases in which a vulnerability or misconfiguration allows an attacker to obtain access to the IdP. This affects all organizations using the identity solution and consequently, they are considered as the worst case scenarios, because of the breadth of the affected audience.
- **IdP Instance and Account Compromise:** The compromise of a single identity provider instance involves a threat model with factors related to misconfigurations, credential leakage, and social engineering attacks against employees. These are threats that target a specific company rather than the IdP solution itself.

The following threat model lays out the various compromises that a company needs to consider depending on the point and severity of attack:

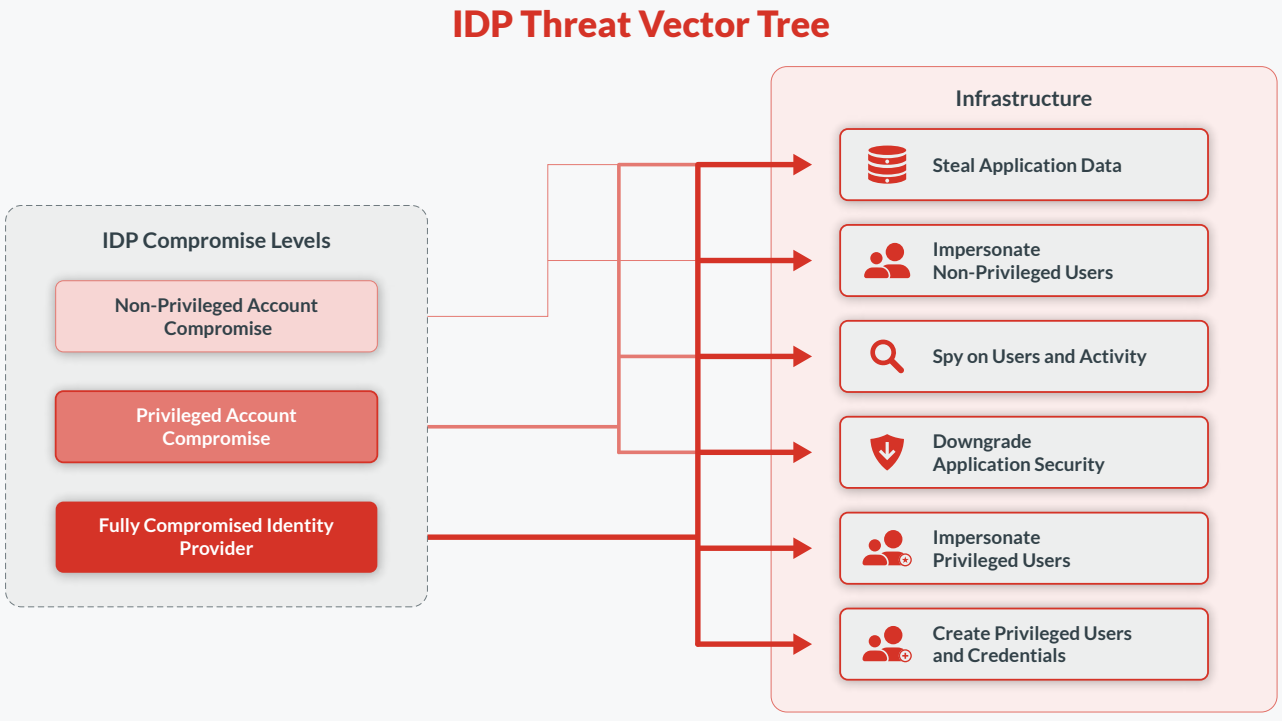

Reducing Impacts from an IdP Compromise

The recent breaches experienced throughout the industry highlight the importance of in-depth security measures to prevent an attacker from compromising an IdP and abusing it to move laterally in the organization. No SSO provider should be assumed to be and remain secure.

> "No SSO provider should be assumed to be and remain secure."
>
> - Luca Carettoni, CEO, Doyensec

The Infrastructure Defense-in-Depth (IDiD) approach consists of implementing an additional layer of security over the identity provider, with protection mechanisms and configurations that defend against the aforementioned threats. The security layer needs to be present on the service providers trusting the IdP as the source of authentication for users.

This approach helps the identity provider and the defense-in-depth provider to work together to strengthen end-user
security. Additionally, the role of the defense-in-depth provider eliminates the need for each Service Provider to independently layer in equivalent levels of defense, which is difficult to achieve with consistency across vendors and is expensive.

Secure Infrastructure Access with Teleport

Teleport, a provider of secure infrastructure access, hardens infrastructure security with on-demand, least privileged access based on cryptographic identity and zero trust, coupled with identity security and policy governance. This architecture eliminates secrets and standing privileges, eliminating broad attack surface leveraged by threat actors:

Modern Access Control	
Governance	Built-in identity and policy governance. Identity locking. Access graph.
Secretless Authentication & Ephemeral Privileges	Least privileged, on-demand access; request workflows. Tamper-proof audit w/ identity-aware session data.
Zero Trust	For applications and workloads. Identity and protocol-aware, reverse tunnels.
Cryptographic Identity	Provisioned and enforced for all users, devices, machines, resources.

Defense-in-Depth Protection Mechanisms

In addition to the architecture described above, companies can employ the following features as protection mechanisms that can be deployed to stop threat actors in compromised IdP scenarios. This chart lays out the role that each key feature can play in blocking threat actors from pivoting from an initial breach to more sensitive data or systems in infrastructure.

Feature	Role in Infrastructure Defense in Depth
Per-session Phishing-Resistant MFA	Protects users against compromises of their infrastructure access, if the threat actor bypasses / controls the MFA on the IdP.
Access Requests	Implements the principle of least privilege, leaving an attacker with no permanent admins to target.
Dual Authorization	Prevents a single successful phishing attack from compromising infrastructure.
Mandatory Phishing-Resistant MFA Enrollment	Prevents weak access patterns.
Web Authentication (WebAuthn)	Secures user identities and prevents attacks like vishing by adding another layer of physical authentication to verify user identities.
MFA for Administrative Actions	Secures against the exploitation of compromised admins, by re-verifying the user's identity promptly before performing any administrative action.
Device Trust	First MFA device enrollment is useful to protect against compromised IdP forcing the auto-provisioning of new users.

Hardening Steps and Best Practices in Configuration

Doyensec further identified in its testing configuration guidelines that can harden security further. These protect against:

- attackers hijacking access reviewer privileges
- roles mapping matching attacks
- exploitation of auto-provisioning of new users, and
- hijacking of both privileged and unprivileged accounts.

Doyensec testing demonstrated that these configurations protected against all tested scenarios of IdP compromise. These recommended configurations are laid out in Appendix B, with a security checklist included in Appendix C.

Why Defense in Depth is Necessary for Modern Infrastructure

By Jack Poller, Principal Analyst, Paradigm Technica

The National Security Agency (NSA) is credited with adapting the military principle of defense in depth to cybersecurity with the goal of providing redundancy in case a security control fails or a vulnerability is exploited. Defense in depth has become the de-facto cybersecurity strategy, enabling organizations to comprehensively address hardware, software, and network security vulnerabilities as well as human factors such as negligence and social engineering.

One of the layers in any cybersecurity stack secures access to modern cloud infrastructure. This layer typically employs one or more identity controls:

- **IAM** – identity and access management, providing account and credential administration, user and device provisioning and deprovisioning, entitlement management, and authentication and single sign on.
- **IGA** – identity governance and administration, which adds processes and policies for segregation of duties, role management, logging, access reviews, analytics, and reporting to IAM.
- **PAM** – privileged access management manages and monitors elevated access and permissions granted to users, accounts, and processes.
- **ITDR** – identity threat detection and response.
- **Passwordless** – replacing passwords with phishing-resistant passwordless authentication.
- **Workload IAM** – IAM for workloads and non-human identities.
- **Authorization** – managing end-to-end authorization of workflows.

Given the massive investment in developing these controls and the plethora of vendors providing these controls, why are they insufficient to protect modern infrastructure?

First, modern infrastructure is vastly different from traditional IT. Whereas in traditional IT compute, storage, and networks are manually pre-provisioned, in modern infrastructure, these resources are automatically and dynamically provisioned on the fly at the same time as applications get deployed using infrastructure as code.

Modern infrastructure environments are predicated on being highly scalable, dynamic, automated, and ephemeral. Thus, they require automated and scalable approaches to security: as more resources are added, as more technology layers are introduced, and as more engineers join the team, the probability of social engineering, human manipulation, or human error leading to credential theft must not increase.

Critically, in modern infrastructure, every access is a privileged access. Gaining access to a DevOps credential gives the attacker carte blanche access to the entire infrastructure, applications, and all the sensitive and critical corporate data. Thus, what works for traditional IT – perimeter security with VPNs, shared secrets, vaults, PAM, IGA, and many other security controls – doesn't work for modern infrastructure.

Second, each of these identity and access controls – separately or in combination – represent a single point of failure. Because these controls are additive and complementary rather than redundant, deploying multiple types of controls doesn't provide an additional layer of security.

For example, in an environment where IAM is deployed, should an attacker use social engineering to compromise a DevOps account, they'll have access to the entire environment. IGA, which provides governance, will only enable the security team to verify that indeed, that account should have privileged access. Likewise, PAM will also enable privileged access because every infrastructure access has to be privileged.

What is needed is to deploy additional security layers beyond the traditional identity and access controls to provide redundancy and eliminate single points of failure for access to critical resources. These additional layers can be classified as "infrastructure defense in depth", or a set of defense in depth layers specifically designed to secure access to modern infrastructure.

Of critical importance is providing redundancy for authentication (AuthN) and authorization (AuthZ). Common practice is to outsource AuthN and AuthZ to a third party identity provider. This means that the security of the AuthN and AuthZ service is outside of the organization's security stack. With infrastructure defense in depth, an additional layer of AuthN and AuthZ is deployed in the organization's stack, providing redundancy and protecting the infrastructure should the IdP be compromised.

Conclusion

Implementing a Defense-in-Depth strategy is essential for businesses looking to defend infrastructure against the compromise of SSO providers and IdPs. By layering multiple security controls, organizations can significantly enhance their ability to detect, prevent, and respond to threats, ensuring a more secure and resilient infrastructure.

Appendix: Infrastructure Defense in Depth with Teleport

A. Protection Mechanisms

The following are examples of protection mechanisms that can be employed in a Teleport cluster, along with the role that they play in stopping threat actors in compromised IdP scenarios.

Feature	What it Is	Role in Infrastructure Defense in Depth
Per-session Phishing-Resistant MFA	Requires additional multi-factor authentication checks when starting a new session for SSH, Kubernetes, databases or desktops.	Protects users against compromises of their infrastructure access, if the threat actor bypasses / controls the MFA on the IdP. This restriction can be enabled resource-wide for a wider audience or per-role to target specific accesses.
Access Requests	Enables users to request access to a resource or role, which can then be approved or denied by a configurable number of approvers. Useful configurations: when a request must be made, what permissions can be requested, how long elevated privileges can last, and how many users can approve or deny different requests.	Implements the principle of least privilege, leaving an attacker with no permanent admins to target. Users receive elevated privileges for a limited period of time. Request approvers can also be configured with limited system access, so they will not be high value targets.
Dual Authorization	Implements Access Requests to require the approval of two team members for a privileged role.	Prevents a single successful phishing attack from compromising infrastructure. With Dual Authorization, Access Requests are restricted to respect certain criteria under the vigilance of multiple reviewers.
Mandatory Phishing-Resistant MFA Enrollment	A mandatory requirement for a user to enroll an MFA device when they create an account, in order to authenticate using that device, when they begin a new session.	Prevents weak access patterns. It should be noted that MFA challenges, at login, apply for local users. SSO users are not impacted by such enforcement.

Web Authentication (WebAuthn)	Supported as a second authentication factor, usable for logging in to the defense-in-depth provider as well as individual SSH nodes or Kubernetes clusters. Supports hardware devices, such as YubiKeys or SoloKeys, and biometric authenticators like Touch ID and Windows Hello.	Secures user identities and prevents attacks like vishing by adding another layer of physical authentication to verify user identities. Even if an attacker successfully tricks a user over the phone, they are still unable to access accounts without the necessary physical keys. By reducing the value and utility of stolen user information, WebAuthn significantly mitigates the risks associated with breaches and the subsequent threats of phishing.
MFA for Administrative Actions	Enforces an additional MFA verification for administrative actions performed in a target application.	Secures against the exploitation of compromised admins, by re-verifying the user's identity promptly before performing any administrative action. This is particularly effective with IdP compromise cases because existing users are no longer usable by attackers to perform administrative actions - de facto blocking many lateral movements.
Device Trust	Enforces the use of trusted devices in addition to the established user's identity and enforced roles, configured as role-based (using RBAC) or cluster-wide.	In particular, first MFA device enrollment is useful to protect against compromised IdP forcing the auto-provisioning of new users. A comprehensive device trust solution supports the following resources: apps (role-based enforcement only), SSH nodes, databases, Kubernetes clusters, and first MFA device enrollment.

B. Hardening Steps and Best Practices in Configuration

This section presents hardening steps and best practices to defend infrastructure against IdP compromise, with Teleport. To develop these recommendations, security researcher Doyensec tested multiple configurations to reproduce threat scenarios. **This research demonstrated that these guidelines protect a Teleport cluster under all tested IdP compromise circumstances.**

Configuration	Doyensec Guidelines
Just-in-Time Access Requests and Dual Authorization Just-in-time Access Requests allow users to request access to a resource or role depending on need, which can then be approved or denied based on a configurable number of approvers. This feature implements the least-privilege principle inside the cluster and controls how privileges are assumed and used. With Dual Authorization, it is possible to enforce the approval of multiple team members for critical actions.	To protect against reviewers that are SSO users with full or limited impersonation capabilities, companies should secure configurations with: • Privileged roles assigned to local users only • SSO roles mappings should follow the ephemeral admin strategy and always require them to request additional privileges • Reviewers that are local users to prevent impersonation and self-acceptance from the IdP (In Teleport, it is not possible to authenticate as an SSO user, with a username which is part of the local users pool)
Authentication Connector and IdP Configuration Security The listed IdP constraints are necessary to prevent automatic provisioning of new Teleport users, forced by generic users or potential Roles Mapping Matching attacks.	To protect against Roles Mapping Matching attacks, companies should use fixed, unique values: **IdP Configuration:** Before setting the authentication connector in Teleport, IdP should be reviewed to identify a username field candidate, which is not editable by end-users and is unique in the IdP's users pool, and a group field candidate, which is editable only by a very restricted group of admin users and is unique in the organization. **Teleport Configuration:** Teleport's secure authentication connector should be configured to prevent automatic inheritance of editor privileges through use of generic names (e.g., admin) or ability to rename groups to acquire editor privileges. Admins should avoid complex string matching in role mapping definitions and instead use fixed values to protect against the exposed attack patterns. Groups usually follow the uniqueness property at the IdP level, hence it will not be possible to have two matches by exploiting a rename operation.

Additional Identity Confirmation Layer via MFA Features

The approach consists of implementing an additional layer of identity confirmation over the identity provider in place.

To mitigate the majority of observed attack patterns, companies should confirm the incoming SSO user's identity by enabling Teleport's various MFA-based features:

- **Per-session MFA** - helps protect session initiations with MFA requirements
- **Administrative Actions MFA Requirement** - prevents most cases of privileges exploitation and escalations from a compromised IdP user
- **WebAuthn As Second Factor** - secures user identities and prevents attacks like vishing, by adding another layer of physical authentication to verify user identities. Reduces the value and utility of stolen user information.
- **Device Trust** - In Teleport v16, Device Trust is extended to the first MFA device enrollment, preventing malicious auto-provisioned SSO users from bypassing MFA-based protections when the trusted device enforcement is configured.

Additional mechanisms and best practices can be found in the Teleport guide: Reducing the Blast Radius of Attacks.

SP-to-IdP Weak Spots

Certain mechanisms in place between the IdP-SP have a significant impact on the overall security. They are not strictly related to Teleport as the service provider and these aspects should be considered in any IdP-SP relationship, where the service provider wants to secure its services, in the event of an IdP compromise.

To protect the SP from being exploited to generate new users:

Auto-Provisioning: Companies can apply design strategies to protect the SP from being exploited by user auto-provisioning, an SSO provider feature that is convenient for user access management at scale but which can be hijacked by threat actors:

- Enforce the usage of trusted devices to enroll the first MFA Device, with device enrollment performed either by a device admin or by the end-user.
- Implement an ephemeral admin strategy, requiring local admin users of the SP to approve or deny Access Requests for privileged actions.
- Protect administrative actions and access to critical resources with a second layer of MFA validation

To protect the SP from being exploited to hijack privileged user accounts:

Access Control Mappings Security: IdPs have different roles and immutable fields usable to create access control mappings. It is important to configure both the IdP and the SP role mappings to minimize the possibility of role / user matching from middle management roles, like team leaders. Using immutable and unique attributes at the IdP to assign roles in the SP prevents any privileged user from manipulating the SSO attributes to escalate privileges in the service provider.

C. Infrastructure Defense in Depth Checklist: Teleport

Doyensec offers the following checklist for customers to verify whether their Teleport cluster has all the available protections and best practices in place to enhance security against IdP compromise scenarios.

- **Just-in-time Access Requests** is configured according to the least-privilege principle; Request reviewers are only local users (i.e., No SSO users as reviewers);
- **Dual-Authorization** is set to further restrict access to administrative actions and implement the concept of ephemeral administrators. Requests reviewers are only local users (i.e., No SSO users as reviewers);
- **SSO Connectors (IdPs)** are configured to restrict roles mappings and automatic provisioning capabilities from non-admin IdP users
  - The username field (IdP-side) mapped as Teleport username is not editable by end-users and is unique in the IdP's users pool;
  - The group field (IdP-side), used to map roles in Teleport, is editable by a very restricted group of users in the IdP and is unique in the organization;
  - The Teleport SSO Connector does not apply lax string matching to map roles. Instead, fixed values from the IdP group are mapped to roles;
- **Device Trust** can be configured to protect against new SSO users being auto-provisioned from a compromised IdP. By enforcing it, new SSO users need to perform the first MFA device enrollment from a trusted device;
- **Access Lists** granting administrative permissions (see RFD 131 [11]) do not have:
  - SSO identities. Only local users should obtain high privileges via access list;
  - Implicit rules referencing attributes obtained from the SSO source;
  - Dangling Identities which are no longer part of the cluster;
- **An additional Identity Confirmation Layer** is applied
  - Per-session MFA is applied cluster-wide to restrict access to various resources with MFA devices;
  - WebAuthn is forced as second factor to avoid OTP-related attacks;
  - Administrative Actions MFA Requirement is active for admin actions, with MFA challenges;
- **Detection & Incident Response Strategies** are in place
  - There are watchdogs listening on the valuable events emitted by Teleport;
  - Moderated Sessions admins are configured as local users, ready to join or assess suspicious sessions;
  - Admins with SSO Connectors management and locking capabilities are ready to be used to block new malicious sessions or invalidate existing ones;
- **Teleport roles do not reference external values** taken from the IdP mappings