Platform Documentation:
API Structure, User Provisioning, and Access Control

1. Introduction

This document provides technical documentation for the Strike48 platform covering API structure, user provisioning, authentication, and access control. It is intended to support technical evaluation and integration planning by engineering and security teams. Strike48 is an API-first, enterprise agentic platform that combines a flexible data access layer with an execution layer composed of automated agents and workflows. The platform enables organizations to query, analyze, and act on operational data while maintaining strong governance over user access and system behavior.

2. Platform and API Architecture

Strike48 exposes its functionality through a unified API layer built on a structured, schemadriven GraphQL interface (https://rc-pov.strike48.com/api/v1alpha/graphiql). This API enables both human users and external systems to interact with platform resources in a consistent and controlled manner.

The platform architecture consists of two primary layers:

- Data layer: Provides access to operational and log-based data across integrated systems without requiring data duplication. The platform supports querying data across existing systems including security tools, observability platforms, and data repositories.
- Execution layer: An agentic framework in which workflows and AI agents perform tasks such as investigation, triage, analysis, and response. Processes can be initiated interactively by users or triggered programmatically via the API.

All interactions are exposed through a single GraphQL API. The schema-driven nature of the API allows consumers to discover available operations and understand required inputs. In appropriate environments, an interactive GraphQL explorer is available for developers to inspect the schema and test queries directly.

3. Authentication

3.1 SAML-Based Single Sign-On

Strike48 supports SAML 2.0 for enterprise authentication and integrates with external identity providers (IdPs) including Okta, Microsoft Entra ID, and other SAML 2.0-compliant providers. Strike48 acts as the Service Provider (SP) in this flow. Strike48 does not store or manage user passwords. All authentication is delegated to the customer's identity provider. The authentication flow is as follows:

- The user navigates to the Strike48 platform and initiates login.
- Strike48 redirects the user to the configured identity provider.
- The identity provider authenticates the user and returns a signed SAML assertion to Strike48.
- Strike48 validates the assertion, establishes a session, and resolves the user's organizational context and role.

Strike48 will work with each customer to configure the SAML integration for their specific identity provider. Required configuration artifacts (Entity ID, ACS URL, SP metadata) are provided by Strike48 during onboarding.

3.2 Session Management

Upon successful SAML authentication, Strike48 issues a session token. All subsequent API requests must include this token. Sessions are scoped to the authenticated user's organizational context and role. Session expiry and timeout policies are configurable at the tenant level during onboarding.

3.3 API Authentication (Machine-to-Machine)

For programmatic and service-to-service access, Strike48 supports API key-based authentication. API keys are issued at the tenant level and scoped to a defined set of permissions.

- API keys are generated and managed through the Strike48 administrative interface.
- Each API key is associated with a defined permission scope and organizational context.
- API keys must be included in the Authorization header of all API requests.
- API keys can be revoked at any time by a tenant administrator.

This mechanism is intended for system integrations, automation pipelines, and agent workflows that operate without a human session context.

4. User Provisioning

4.1 Provisioning Models

Strike48 supports two user provisioning models, which can be used independently or in combination:

Model	Description	When to Use
Just-in-Time (JIT)	User account is automatically created in Strike48 on the user's first successful SAML login. The account is provisioned with a default role that can be updated by an administrator.	Preferred for most deployments Reduces administrative overhead and keeps lifecycle aligned with the IdP.
Pre-Provisioned	User accounts are created in Strike48 by an administrator before first login, with role and context assigned in advance. The user's SAML identity is linked on first login.	Used when roles must be assigned before a user's first access, or when tighter preapproval controls are required.

4.2 User Deprovisioning

Because authentication is managed through the customer's identity provider, disabling or removing a user in the IdP immediately prevents that user from authenticating to Strike48. Active sessions are invalidated upon the next request validation cycle.

Administrators can also explicitly deactivate user accounts within Strike48 through the administrative interface or via the API, independently of IdP state.

4.3 User Lifecycle via API

User lifecycle operations are available programmatically through the Strike48 GraphQL API. Authenticated administrators can retrieve, create, update, and deactivate users within their organizational scope.

Example — retrieve users within the current organizational context:

```
query {
  users {
    id
    email
    role
    status
  }
}
```

All user API operations are subject to the caller's role-based permissions and organizational scope. A user cannot retrieve or modify accounts outside their organizational context.

5. Role-Based Access Control (RBAC)

5.1 Overview

Strike48 employs a role-based access control (RBAC) model to manage permissions across the platform. Every user is assigned a role that governs the operations they may perform and the data they may access.

Role definitions — including the specific permissions associated with each role — are configured during tenant onboarding in collaboration with Strike48. The role model is designed to accommodate the customer's existing organizational structure and access governance requirements.

5.2 Permission Enforcement

When a user issues a request, the platform evaluates that request against the user's assigned role. Requests that exceed the user's permissions are rejected. This evaluation is applied consistently across all API interactions, regardless of whether the request originates from a human user or an automated process.

Permissions are enforced at multiple levels:

- Operation level: Whether the user may invoke a given API operation (e.g., execute a workflow, manage users, view a dashboard).
- Data level: Which records and fields the user may read or modify within an operation.
- Organizational scope: Users can only access data and resources belonging to their assigned organizational context.

5.3 Fine-Grained Authorization

In addition to operation-level permissions, Strike48 enforces authorization at the level of individual data elements. A user with permission to retrieve a set of records may not be authorized to view all fields associated with those records. In such cases, only permitted fields are returned; unauthorized fields are omitted from the response.

This model extends to the execution layer. Automated agents and workflows operate within defined permission scopes, with access restricted to specific data sources and tools. These scoped controls ensure that automated processes execute only within their authorized boundaries and cannot escalate privileges beyond what the initiating context permits.

6. Organizational Scope and Data Isolation

Strike48 is designed to support multi-tenant environments where multiple organizations operate within the same platform infrastructure. The platform enforces strict data isolation between tenants.

Each user is associated with a specific organizational context. All API requests are evaluated within that context. A user cannot access data belonging to a different organization, even if they share an identical role with a user in another tenant. Organizational context is derived from the authenticated identity and cannot be overridden by request parameters.

For organizations that operate hierarchical structures — such as enterprises with distinct business units or managed service environments — Strike48 supports a multi-tier domain scoping model. Access rights and data visibility can be configured at each tier, enabling parent organizations to maintain administrative oversight while preserving isolation between child contexts. This capability is configured during onboarding.

7. Audit Logging

Strike48 maintains an audit record of user actions and API operations within the platform. Audit events are surfaced in the Strike48 administrative interface and are accessible to authorized administrators within the tenant.

Logged events include:

- User authentication events (login, logout, failed attempts)
- User provisioning and deprovisioning actions
- Role assignment and modification
- API key creation and revocation
- Workflow and agent execution events
- Data access operations by privileged users

Audit log retention and export requirements are discussed during onboarding. Customers with requirements to forward audit events to external systems should raise this during the onboarding process.

8. Security Principles

Access control in Strike48 is governed by the following principles:

- Least privilege: Users and automated processes are granted only the permissions necessary to perform their defined functions.
- Default deny: Access is denied unless explicitly permitted by the user's assigned role.
- Consistent enforcement: Authorization checks are applied uniformly across all API interactions. There are no alternative access paths that bypass permission evaluation.
- Separation of identity: Human user sessions and API key-based service identities are managed and audited independently.
- Tenant isolation: Organizational data boundaries are enforced at the platform level and cannot be overridden by application-layer requests.

9. Summary

Strike48 provides a secure, enterprise-grade platform for agentic data access and workflow execution. Its authentication model is built on SAML 2.0, delegating credential management to the customer's existing identity provider and supporting both JIT and pre-provisioned user lifecycle models.

Permissions are enforced through a role-based access control model applied consistently across all API interactions, extended by fine-grained field-level authorization and scoped execution controls for automated agents. Multi-tenant data isolation ensures that organizational boundaries are strictly maintained at the platform level.

Audit logging provides visibility into user and system activity, accessible through the Strike48 administrative interface.

Specific role definitions, SAML configuration details, domain hierarchy structure, and audit export requirements are established collaboratively during the tenant onboarding process. Customers with questions not addressed in this document should direct them to their Strike48 technical contact.