



ManageEngine SecureAI

Governance and runtime protection
for agentic AI

AI has stopped answering questions. It now acts.

01 Agents read from and write to real systems: ticketing, CRM, cloud, code repositories

02 They call tools through MCP (Model Context Protocol) servers, often with standing credentials

03 They're being spun up across Copilot Studio, AI Foundry, Gemini Enterprise Agent Platform, Zia, and more—frequently without IT in the loop

Agents are everywhere. Governance isn't.

Every agent is a new kind of identity, and most are ungoverned:

- **Unmanaged:** Discovered, but with no responsible human. No one accountable when something goes wrong
- **Over-privileged:** Able to access, modify, or exfiltrate sensitive data
- **Unmonitored at runtime:** Tool calls happen with no inspection or policy

Traditional IAM wasn't built for this


- Built for humans and service accounts, not autonomous agents that reason and chain tool calls
- No native concept of an agent's model, instructions, knowledge sources, or tools
- No visibility into agent-to-tool traffic (the MCP layer)
- Compliance frameworks now expect AI-specific controls that legacy IAM was never designed to evidence

Meet ManageEngine SecureAI

A single platform to govern and protect agentic AI across every major platform.
SecureAI lets an organization:

Capability	What it does
Discover	What agents, models, and MCP servers exist?
Govern	Who owns each agent, and what access is justified?
Protect	What can an agent actually do when it calls a tool?
Assess	How risky is our AI estate, and can we prove compliance?
Operate	Who administers SecureAI, and what's happening across the estate?

One pane of glass across the AI agents of Microsoft, Google, and Zoho—plus risk assessment for the enterprise AI platforms your teams already use.



Every major AI platform, governed from one console

SecureAI connects to two kinds of platform:

Agent platforms: Discover, govern, and protect at runtime



AI Foundry



Copilot Studio



Gemini Enterprise
Agent Platform



Zia Agents

Enterprise AI platforms: Risk assessment



Anthropic Claude



GitHub Copilot



Gemini

Identity providers




Entra ID

Google Workspace

Google Workspace

Bring every agent in

Configure the platforms you want to govern. Each connection uses scoped authorization, so SecureAI reads only what it needs to discover and manage agents.

The screenshot shows the 'Instance Creation - Microsoft' dialog box. It has a close button (X) in the top right corner. The dialog is divided into two tabs: 'General' (selected) and 'Connection'. The 'General' tab contains the following sections: 'General' with an 'Instance Name' field (containing 'Microsoft Instance 1') and a 'Description' link; 'Choose Capabilities' (empty); 'Identity Provider & Protection' with an unchecked checkbox for 'Entra ID' (with a description: 'Acts as the identity provider for secure access protection.'); and 'Governance' with two unchecked checkboxes: 'AI Foundry' (description: 'Use as a data source to discover and manage agents deployed within your AI Foundry environment.') and 'Copilot Studio' (description: 'Use as a data source to discover and manage agents built and deployed through Copilot Studio.'). At the bottom are 'Next' and 'Cancel' buttons.

Connected sources, always in view

Each connected stays live in the console. View instance name, description, and connected services.

[Dashboard](#) [Governance](#) [Access Protection](#) [Reports](#) [Delegation](#) [Settings](#) [Support](#) [Application Settings](#)

Application Settings

Here you can manage your Directory Identity sources. [Learn More](#)

Manage Applications


1 - 13 of 13

Actions	Instance Name	Description	Application Name	Last Fetch Time	Fetch Interval	Associated Services	Status
 	secureai.com	-	Microsoft	9/2/2026, 2:21:35 PM Fetch	Every 6 hours 	Entra ID, Copilot Studio	Sync failed
 	secureai.com	-	Anthropic Claude	-	-	Risk Assessment	Configured
 	Github Copilot - secureai.com	-	Github Copilot	-	-	Risk Assessment	Configured
 	Zia Instance 1	-	Zia	9/2/2026, 5:39:28 PM Fetch	Every 6 hours 	Zia Agents	Sync completed
 	Google Instance 1	-	Google	9/2/2026, 2:44:54 PM Fetch	Every 6 hours 	Gemini Enterprise Agent Platform, Go...	Sync failed
 	Microsoft Instance 122	-	Microsoft	-	-	Entra ID, AI Foundry, Copilot Studio	Not Configured
 	Anthropic Claude Instance33	-	Anthropic Claude	-	-	Risk Assessment	Not Configured
 	Microsoft Instance 1gqsd	-	Microsoft	-	-	Entra ID, AI Foundry	Not Configured
 	Microsoft Instance 1233	-	Microsoft	-	-	Entra ID, AI Foundry	Not Configured
 	Microsoft Instance 8	-	Microsoft	-	-	Entra ID, Copilot Studio	Not Configured
 	Anthropic Claude Instance 1	-	Anthropic Claude	-	-	Risk Assessment	Not Configured

Your AI posture at a glance

The dashboard surfaces the numbers a security lead actually asks for, including:


- Total agents and MCP servers
- Unmanaged and privileged agents
- Agents without guardrails
- An enterprise AI risk score (posture indicator)
- Plus agents by platform, top MCP servers, guardrail coverage, and full AI inventory.



Every agent, accounted for

See every discovered agent in one list, including unmanaged and high-privileged agents.
For any agent, generate a visual map of everything it can reach:

- Agent identity and permissions (API permissions, Entra roles, Azure RBAC roles)
- Knowledge sources and tools
- Connected MCP servers and the model behind it
- Guardrails configured on the agent



What your agent runs on: MCP servers and models

MCP servers: Every server in use, with its endpoint URL, application, instance, which agents use it, and its protection status. Unprotected servers can be secured by registering them and binding them to a gateway.

Models: A full inventory of models in play: name, model ID, provider, application, and which agents use each.

Dashboard	Governance	Access Protection	Reports	Delegation	Settings	Support	Application Settings
▼ AI Inventory							
AI Agents							
MCP Servers							
Models							
▼ Access Configuration							
Ownership Policy							
Entitlement							
Access Templates							
▼ Access Certification							
Access Review Request							
Access Certification Campaign							
Governance Assigning Rule							
▼ AI Risk Assessment							
AI Risk Assessment							

MCP Servers

Lists all discovered MCP servers, their endpoints, the agents using them, and their protection status.

1 - 24 of 24

MCP Server	Endpoint URL	Application	Instance	Used By	Protection
new_atlassian	https://mcp.atlassian.com/v1/sse	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_azure	https://azuremcp.azure.com/mcp	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_box	https://mcp.box.com/mcp	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_cliq_zxo	https://cliq-messaging-60068503767.zohomcp.in/mcp/04d34d716aa1940de139ee901f5978d/message	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_dropbox	https://mcp.dropbox.com/mcp	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_gate	https://7fdb-182-74-243-51.ngrok-free.app/api/mcp-proxy	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_github	https://api.githubcopilot.com/mcp/	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_gmail	https://workspace.google.com/mcp/gmail	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_linear	https://mcp.linear.app/mcp	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_manoj-zoho-mcp	https://zohowriter-60081830398.zohomcp.in/mcp/e9cbc158cc0b7a9c16ee2b8edeae80db/message	Copilot Studio	secureai.com	Revenue Growth Analyst	⊗ Unprotected Configure
new_mcp cliq	https://cliq-messaging-60068503767.zohomcp.in/mcp/775dab8a342e4121005d6867d25b8c07/message	Copilot Studio	secureai.com	-	⊗ Unprotected Configure
new_mslearn	https://learn.microsoft.com/api/mcp	Copilot Studio	secureai.com	-	⊗ Unprotected Configure

Knowing what exists isn't enough

Governance decides what an agent should be able to do. At runtime, an agent makes live tool calls to real systems, and that's where the risk is realized:

Without inspection, a prompt-injected or misbehaving agent can issue destructive or sensitive calls

Standing credentials mean an agent often has more reach than any single task requires



SecureAI inserts an MCP gateway between agents and tools, so every call is identity-bound and policy-checked. This is what legacy IAM cannot do.

One gateway, every call inspected

Three steps to bring a tool under control:

- **Register the MCP server:** Add the upstream URL and authentication to the registry
- **Create a gateway:** Group servers under one signed URL with shared identity and policy rules
- **Connect the client:** Point the agent (Copilot Studio, AI Foundry agent, custom code) at the gateway URL


Each gateway exposes a single URL where every request is inspected—federating the real MCP servers behind one identity contract, with conditional access, tool-level access control, DLP, and logging applied centrally.

One gateway, every call inspected



Dashboard	Governance	Access Protection	Reports	Delegation	Settings
▼ Gateway					
MCP Server Registry					
MCP Gateways					
▼ Access Controls					
Control Policy					
▼ Monitoring					
Gateway Activity					
Insight Reports					

MCP Server Registry		
MCP servers registered with the gateway. Tool calls are evaluated against access		
🔍	🔼	🟢 Enable 🔴 Disable 🗑️ Delete
☐	Actions	MCP Server Name
☐		Jira MCP
☐		Notion MCP
☐		Sentry MCP
☐		Zapier MCP

MCP Gateway Onboarding Steps



One gateway, every call inspected



Block the dangerous call, allow the rest

Control policies evaluate every tool call at runtime against context (IP, time, user attributes, MFA, agent risk score). Two policy types: Access Policy allows or blocks the call, and DLP Policy redacts or blocks sensitive data in it.

Examples:

- Access Policy
 - Block destructive actions outside business hours
 - Block bulk or wildcard delete operations
 - Restrict a tool to a named set of agent identities
- DLP Policy
 - Redact sensitive fields before they leave the response
 - Block calls carrying regulated data to unapproved servers

Rules build from context—time, IP, location, tool, role, and data type.

Block the dangerous call, allow the rest

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Application Settings

Gateway

MCP Server Registry

MCP Gateways

Access Controls

Control Policy

Monitoring

Gateway Activity

Insight Reports

Control Policy

Context-aware rules enforced at runtime @ IP, time, user attributes, MFA, agent risk score. Each policy is attached to one or more MCP Gateways.

All OWASP Top 10 MITRE ATLAS

Policy Parameters

+ New Control Policy

1 - 25 of 66

	Actions	Policy Name	Description	Compliance	Policy Type	Action	Status
☐		AfterHours-Destructive-Plus-Injection	Blocks injection payloads sent to destructive tools outside business hours.	ATLAS - AML.T0053 — AI Agent Tool Invocation +1	DLP Policy	Block	enabled
☐		Block-Adversarial-Suffix-Patterns	Blocks adversarial suffix patterns in tool input.	ATLAS - AML.T0054 — LLM Jailbreak +1	DLP Policy	Block	enabled
☐		Block-Adversarial-Write-Payloads	Blocks adversarial payloads sent to write tools.	ATLAS - AML.T0020 — Poison Training Data +1	DLP Policy	Block	enabled
☐		Block-AfterHours-Destructive-Writes	Blocks destructive write operations outside business hours.	ATLAS - AML.T0034 — Cost Harvesting +1	Access Policy	Block	enabled
☐		Block-Agent-Self-Modification	Blocks tool calls that attempt agent self-modification.	ATLAS - AML.T0110 — AI Agent Tool Poisoning +2	DLP Policy	Block	enabled
☐		Block-Biometric-Patterns	Blocks biometric identifiers in tool input and responses.	ATLAS - AML.T0024 — Exfiltration via AI Inference API	DLP Policy	Block	enabled
☐		Block-Bulk-Operations-In-ToolParam	Blocks bulk-operation requests in tool input parameters.	ATLAS - AML.T0053 — AI Agent Tool Invocation +1	DLP Policy	Block	enabled
☐		Block-CRLF-Injection	Blocks CRLF-injection sequences in tool input.	ATLAS - AML.T0050 — Command and Scripting Interpreter +1	DLP Policy	Block	enabled
☐		Block-Dark-Patterns-Outbound	Blocks dark-pattern content sent through outbound tools.	OWASP - ASIO6 — Human Agent Trust Exploitation	DLP Policy	Block	enabled
☐		Block-Data-Residency-Violation	Blocks tool responses that would violate data-residency rules.	ATLAS - AML.T0024 — Exfiltration via AI Inference API	DLP Policy	Block	enabled

Block the dangerous call, allow the rest

Dashboard Governance Access Protection Reports Delegation Settings

Gateway

- MCP Server Registry
- MCP Gateways

Access Controls

- Control Policy

Monitoring

- Gateway Activity
- Insight Reports

Control Policy
Context-aware rules enforced at runtime @ IP, time, user attributes, M

All OWASP Top 10 MITRE ATLAS

Enable Disable Delete

	Actions	Policy Name
☐		AfterHours-Destructive-Plus
☐		Block-Adversarial-Suffix-Patt
☐		Block-Adversarial-Write-Payl
☐		Block-AfterHours-Destructiv
☐		Block-Agent-Self-Modificati
☐		Block-Biometric-Patterns
☐		Block-Bulk-Operations-In-To
☐		Block-CRLF-Injection
☐		Block-Dark-Patterns-Outbou
☐		Block-Data-Residency-Violat

New Control Policy

General

Policy Name * [Description](#)

Policy Type

Identities to apply

Conditions

Conditions

① No conditions specified
[Click here](#) to add Condition


Policy Action

Action ☒ Allow & Audit ☐ Block

From what exists to what's happening

The Access Protection view shifts from "what exists" to "what's happening" at runtime:

- Active vs. inactive MCP servers
- Live tool-call outcomes (allowed vs. blocked)
- Access and DLP policy activity
- Gateway latency percentiles



MCP Server Registry							
MCP servers registered with the gateway. Tool calls are evaluated against access policies before forwarding. Learn more							
🔍 Enable Disable Delete Showing All 1-25 of 100							
Actions	MCP Server Name	Description	Authentication Type	Available Tools	Server Status		
☒	Atlassian Sever	Atlassian Sever	Per-user authentication	getTicket, updateTicket more	Active		
☐	ServiceNow Server	ServiceNow Server	Shared credentials	getTicket, updateTicket more	Inactive Connect		
☐	Zoho Server	Zoho Server	Shared credentials	getRecord, updateRecord more	Active		

A receipt for every action

Gateway activity: A real-time log of every MCP call: timestamp, agent, identity, tool, latency, policy action, policies matched, status. Drill into any request to see exact arguments and response.

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Application Settings

Gateway

MCP Server Registry

MCP Gateways

Access Controls

Control Policy

Monitoring

Gateway Activity

Insight Reports

Gateway Activity

MCP gateway activity recorded by the gateway. Tool calls are listed newest first.

Filter

1 - 25 of 1000

Timestamp	Agent Principal Name	Tool Name	Policy Action	Status	Latency (ms)	
2026/08/30 12:05:51	priya.nair@corp.local	filesystem.read_file	allow	Success	149	Details
2026/08/30 12:05:28	alice.martin@corp.local	-	deny	Failure	2340	Details
2026/08/30 11:56:43	dev.svc@corp.local	-	allow	Success	58	Details
2026/08/30 11:44:01	bob.chen@corp.local	-	allow	Success	169	Details
2026/08/30 11:20:46	ci-runner@corp.local	-	allow	Failure	2927	Details
2026/08/30 11:10:29	alice.martin@corp.local	jira.add_comment	deny	Failure	4788	Details
2026/08/30 11:02:23	dev.svc@corp.local	-	allow	Success	19	Details
2026/08/30 10:46:17	bob.chen@corp.local	postgres.list_schemas	allow	Success	25	Details
2026/08/30 10:32:09	carol.diaz@corp.local	slack.search_messages	allow	Success	170	Details
2026/08/30 10:25:49	bob.chen@corp.local	github.list_pull_requests	allow	Success	90	Details
2026/08/30 10:17:29	dev.svc@corp.local	-	allow	Success	103	Details
2026/08/30 10:02:42	bob.chen@corp.local	jira.search_issues	allow	Failure	4311	Details

A receipt for every action

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Gateway

MCP Server Registry

MCP Gateways

Access Controls

Control Policy

Monitoring

Gateway Activity

Insight Reports

Gateway Activity

MCP gateway activity recorded by the gateway. Tool calls are listed newest first.

Timestamp	Agent Principal Name	Tool Name	Policy Action
2026/08/30 12:05:51	priya.nair@corp.local	filesystem.read_file	allow
2026/08/30 12:05:28	alice.martin@corp.local	-	deny
2026/08/30 11:56:43	dev.svc@corp.local	-	allow
2026/08/30 11:44:01	bob.chen@corp.local	-	allow
2026/08/30 11:20:46	ci-runner@corp.local	-	allow
2026/08/30 11:10:29	alice.martin@corp.local	jira.add_comment	deny
2026/08/30 11:02:23	dev.svc@corp.local	-	allow
2026/08/30 10:46:17	bob.chen@corp.local	postgres.list_schemas	allow
2026/08/30 10:32:09	carol.diaz@corp.local	slack.search_messages	allow
2026/08/30 10:25:49	bob.chen@corp.local	github.list_pull_requests	allow
2026/08/30 10:17:29	dev.svc@corp.local	-	allow
2026/08/30 10:02:42	bob.chen@corp.local	jira.search_issues	allow

Request Details

Timestamp : 2026/08/30 12:05:51

Agent Name : claude-desktop

Identity : priya.nair@corp.local

Tool : filesystem.read_file

MCP Server Name : srv-fs-01

Latency : 149 ms

Policy Action : allow

Policies Matched : filesystem-access-policy [more](#)

Status : Success

Status Message : Tool call completed successfully

Arguments

{
 "path": "/srv/app/config/settings.yaml"
}

Response

{
 "result": {
 "content": {
 {
 "text": "ok",
 "type": "text"
 }
 }
 },
 "jsonrpc": "2.0"
}

Close

Insight reports

Agent principal usage, MCP tools usage, MCP server latency, inactive MCP servers, MCP server errors, control policy audits.

The screenshot displays the ManageEngine Insight Reports dashboard. The top navigation bar includes links for Dashboard, Governance, Access Protection, Reports, Delegation, Settings, and Support. The 'Reports' tab is currently selected. On the left, a sidebar menu lists categories: Gateway (with sub-items MCP Server Registry and MCP Gateways), Access Controls (with sub-item Control Policy), and Monitoring (with sub-items Gateway Activity and Insight Reports, which is highlighted). The main content area is titled 'Insight Reports' and contains a descriptive subtitle: 'Prebuilt SecureAI Gateway reports â€” agent usage, tool activity, errors, latency, and policy audits.' Below this, the reports are organized into three columns: 'Agent Principal Reports' (containing 'Agent Principal Usage Report'), 'MCP Based Reports' (containing 'MCP Tools Usage Report', 'MCP Servers Latency Report', 'MCP Server Error Reports', and 'Inactive MCP servers'), and 'Control Policy' (containing 'Control Policy Audit'). An 'Application Settings' button is located in the top right corner of the dashboard area.

Agent Principal Reports	MCP Based Reports	Control Policy
Agent Principal Usage Report	MCP Tools Usage Report	Control Policy Audit
	MCP Servers Latency Report	
	MCP Server Error Reports	
	Inactive MCP servers	

The agents that need watching

Unmanaged agents: Agents with no responsible human; assign an owner directly from the list.

High-privileged agents: Agents with elevated permissions that can access, modify, or exfiltrate sensitive resources; flagged for closer oversight and stricter controls. The privileged list can be added to and modified.

The screenshot displays the ManageEngine AI Agents interface. The top navigation bar includes links for Dashboard, Governance, Access Protection, Reports, Delegation, Settings, and Support. The left sidebar shows a tree view with categories like AI Inventory, Access Configuration, and Access Certification. The main content area is titled 'AI Agents' and shows a list of unmanaged agents. The table has columns for Actions, Agent Name, Application, Model, Last Updated, and Owners. Each row includes an 'Assign Owner' link.

Actions	Agent Name ↑	Application	Model	Last Updated	Owners ⓘ
☐	Sales Agent	AI Foundry	GPT-4	2025-12-20 01:29:58	Assign Owner
☐	AI Assistant	Vertex AI	GPT-4	2025-12-20 01:29:58	Assign Owner
☐	M365 - Strategist	AI Foundry	GPT-4	2025-12-20 01:29:58	Assign Owner
☐	AWS Agent	Bedrock	GPT-4-turbo	2025-12-20 01:29:58	Assign Owner
☐	Zia Help Agent	Zia Agents	Opus4.6	2025-12-20 01:29:58	Assign Owner

Every agent is accountable: Ownership policy

- Every agent and agent identity needs a human owner held accountable
- Policies auto-assign owners to newly discovered agents based on attribute conditions, first match by priority
- Manually assigned owners are preserved—re-running a policy never overwrites a hand-set owner
- Assignment history is tracked; configuration is product-level and prioritize-able

Dashboard Governance Access Protection Reports Delegation Settings Support

AI Inventory

- AI Agents
- MCP Servers
- Models

Access Configuration

- Ownership Policy
- Entitlement
- Access Templates

Access Certification

- Access Review Request
- Access Certification Campaign
- Governance Assigning Rule

AI Risk Assessment


- AI Risk Assessment

Ownership Policy

Auto-invoked when new agents are discovered.

🔍 | ☒ Enable ☐ Disable ☐ Delete

☐	Actions	Name
☐	☒ <input type="edit"/> <input type="trash"/>	Default Ownership Policy
☐	☒ <input type="edit"/> <input type="trash"/>	Tech Admin Policy
☐	☐ <input type="edit"/> <input type="trash"/>	Fallback Policy



The right access, every time

- **Entitlements:** Fetched from every configured agent platform, with name, description, type, application, and owner. Reassign owners individually or in bulk.
- **Access templates:** Bundle entitlements into reusable definitions spanning identity (Entra roles, API permissions, Azure RBAC roles) and agent configuration (model, alternate model, MCP servers, knowledge) — used to scope exactly what a certification campaign reviews.

The right access, every time

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Application Settings

AI Inventory

AI Agents

MCP Servers

Models

Access Configuration

Ownership Policy

Entitlement

Access Templates

Access Certification

Access Review Request

Access Certification Campaign

Governance Assigning Rule

AI Risk Assessment

AI Risk Assessment

Entitlements

Lists the entitlements from the directories/applications and their associated JIT policies, along with an assignment summary. [Learn more](#)

Select Instances

All Selected

🔍

👤 Update Owner

1 - 25 of More < > 📄 ↺

☐	Actions	Name	Description	Type	Directories	Owners ⓘ
☐	🔗	Attribute Log Reader	Read audit logs related to custom security attributes.	Role	🔗 Microsoft Instance 1	-
☐	🔗	Compliance Data Administrator	Creates and manages compliance content.	Role	🔗 Microsoft Instance 1	-
☐	🔗	Azure Information Protection Administrator	Can manage all aspects of the Azure Information Pro...	Role	🔗 Microsoft Instance 1	-
☐	🔗	Attack Simulation Administrator	Can create and manage all aspects of attack simulati...	Role	🔗 Microsoft Instance 1	-
☐	🔗	Extended Directory User Administrator	Manage all aspects of external user profiles in the ex...	Role	🔗 Microsoft Instance 1	-
☐	🔗	Read all assignments	Allows the app to read data for all assignments in the...	Api Permission	🔗 Microsoft Instance 1	-
☐	🔗	Read all pending external user profiles	Allows the app to read available properties of pendi...	Api Permission	🔗 Microsoft Instance 1	-
☐	🔗	Organizational Messages Writer	Write, publish, manage, and review the organization...	Role	🔗 Microsoft Instance 1	-
☐	🔗	Read and write all on-premises directory synchronization information	Allows the app to read and write all on-premises dire...	Api Permission	🔗 Microsoft Instance 1	-
☐	🔗	Viva Pulse Administrator	Can manage all settings for Microsoft Viva Pulse app	Role	🔗 Microsoft Instance 1	-
☐	🔗	Read and write all self-initiated courses	Allows the app to create, update, read and delete all ...	Api Permission	🔗 Microsoft Instance 1	-

The right access, every time

[Dashboard](#) [Governance](#) [Access Protection](#) [Reports](#) [Delegation](#) [Settings](#) [Support](#) [Application Settings](#)

←

Create Access Template

The access template bundles entitlements from various directories and applications. [Learn more...](#)

* Template Name

[Description](#)

* Select Instances

Microsoft Instance 122, secureai.com

+

Category: Default

Enable Drag-and-drop

Microsoft Instance 122

secureai.com

Google Instance 1

Copilot Entitlements

Entra Roles

+

API Permissions

+

MCP Servers

+

Knowledge

+

AI Foundry Entitlements

Entra Roles

+

API Permissions

+

Azure RBAC Roles

+

Model

+

MCP Servers

+

Knowledge

+

Save Template

Cancel

Keep access honest

A per-platform risk view across Anthropic Claude, Gemini, and GitHub Copilot. Each account gets:

- A risk score based on detected misconfigurations (score shown in UI is illustrative)
- Findings grouped by severity across categories
- Concrete remediation guidance for each finding

Risk assessment covers these platforms whether or not their agents are managed in SecureAI.

Keep access honest

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Application Settings

AI Inventory

AI Agents

MCP Servers

Models

Access Configuration

Ownership Policy

Entitlement

Access Templates

Access Certification

Access Review Request

Access Certification Campaign

Governance Assigning Rule

AI Risk Assessment

AI Risk Assessment

Access Certification Campaign

Create access audit campaigns in order to review the access of users or groups to resources and assign certifiers to approve or revoke access as required. [Learn more...](#)

1 Campaign Details

2 Certifier & Scheduler

3 Settings

4 Summary

General Details

Provide a name, description and set priority for the campaign.

Campaign Name *

Description

Priority

Normal

* Entitlement Selection

Select the entitlements that should be reviewed in this campaign.

Select Entitlements / Access Templates *

- Select -

+

Object Selection

Select the users or groups to be reviewed. By default, all the users or groups that are assigned the selected entitlement are included.

Object Selection Type

All Agents

Cancel

Step 1 of 4

Next

Right request, right approver

Governance assigning rules route each request to the right technicians based on criteria, in priority order. Requests flow through a configurable business workflow:

Requester → Reviewer → Approver → Executor

The number of reviewers/approvers can be configured per stage.

The image displays the ManageEngine Governance Assigning Rule configuration interface and a modal dialog for creating a business workflow.

Background Interface (Governance Assigning Rule):

- Navigation:** Dashboard, Governance, Access Protection, Reports, Delegation, Settings, Support.
- Left Sidebar:**
 - AI Inventory
 - AI Agents
 - MCP Servers
 - Models
 - Access Configuration
 - Ownership Policy
 - Entitlement
 - Access Templates
 - Access Certification
 - Access Review Request
 - Access Certification Campaign
 - Governance Assigning Rule**
 - AI Risk Assessment
 - AI Risk Assessment
- Main Content Area:**
 - Governance Assigning Rule**
Create rule to assign requests to appropriate technicians.
 - * Rule Name:** [Text Field]
 - Description:** [Text Field]
 - Business Workflow:** Default certifier business workflow (Dropdown)
 - Request Administration:**
 - Request Creation
 - Request Execution
 - Request Expiration
 - Buttons:** Add Rule, Cancel

Modal Dialog (Create Business Workflow):


- Workflow Name:** [Text Field]
- Description:** [Text Field]
- Workflow Stages:**
 - Requester:** The one who raises a request for a particular action. [Continue]
 - Reviewer:** The one who assesses the request, weighs its pros and cons, and offers recommendations. [Continue]
No. of Reviewers: 1
 - Approver:** The one who possesses the authority to finalize an action. [Continue]
No. of Approvers: 1
 - Executor:** The one who executes the approved action. [Continue]
- Buttons:** Create Workflow, Cancel

Know your risk, by platform

A per-platform risk view across Anthropic Claude, Gemini, and GitHub Copilot. Each account gets:

- A risk score and a score trend over time (score shown in UI is illustrative)
- Findings grouped by severity across categories
- Concrete remediation guidance for each finding
- Compliance mapping to the EU AI Act, ISO/IEC 42001, NIST AI RMF, Gartner AI TRiSM, and NCSC

Know your risk, by platform



Each finding comes with a fix

Every finding comes with evidence and step-by-step remediation—so a finding becomes an action, not just an alert.

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Application Settings

Identity Risk Assessment

API Key Security

API keys with unclear names

Inactive API keys

API keys in the Default workspace

Service account API keys require separate governance

Service account keys exceed the approved rotation interval

Workspaces Above the Active Key Limit

API keys without a documented owner

API keys exceed the approved rotation interval

Data Governance

Identity & Access Management

Monitoring & Logging

Network Access Control

Third-Party Integrations

← API keys in the Default workspace

42%

Full lifecycle control of API keys and service-account credentials, including complete inventory, maximum exist. It disables unused or inactive keys, enforces descriptive naming, and tracks service-account-owned

Severity : HIGH | Weight : 8 | Status : Fail

Risk Indicators - 5 of 12

Anthropic Claude: secureai.com

Name	ID	Workspace
no use	apikey_012j2QqtXZiqY9kTBpKw5b9	(none)
dumm4	apikey_0181tibvt3nvdnesyfoEAef	(none)
svac_018BXZX3hKtIiHJatoDymyif	apikey_013ap2HjoC2xbhmEjtp2Hmre	(none)
d	apikey_01LNUWUPQ7d94LkvevNsvu2LU	(none)
dummy3	apikey_014mTMyBfVZ4V1igDGsHogEr	(none)

API keys in the Default workspace

Description

Bind each API key to a specific workspace to limit access and improve usage and cost attribution.

Remediation measures

1. Confine each credential to one workload. A key in the Default workspace shares that workspace's limits and attribution. Use a dedicated workspace when the workload requires separate access and cost boundaries.
2. Go to Claude Console (console.anthropic.com) > Settings > API Keys / Workspaces.
3. Find the keys with no workspace of their own, shown against Default.
4. Create equivalent keys inside the correct workspace and repoint consumers.
5. Revoke the Default-workspace keys.
6. Use separate workspaces to isolate development, staging and production.
7. If leaving these keys unscoped is a deliberate, documented decision, record an exception through your governance process.

Compliance Frameworks (Domain Level)

- EU AI Act Art.15
- NIST AI RMF MANAGE
- OWASP LLM02:2025
- ISO 42001 A.5
- MITRE ATLAS Credential Access
- Gartner AI TRISM
- NCSC Secure-deployment
- NIST SP 800-53 IA-5/AC-2
- CIS v8 5

Delegate without over-granting

Hand off day-to-day SecureAI administration without handing out full control.

- **Help desk roles** define granular permissions across Governance, Access Protection, Reports, and Administration—down to specific inventory, access-management, and certification actions.
- **Help desk technicians** are assigned roles, tenants, and resources, plus the platforms they can act on.
- **Audit and admin audit reports** record who did what.
- **Logon settings** (CAPTCHA, lockout, SSO, 2FA, IP restrictions) secure technician access.

Delegate without over-granting

Dashboard Governance Access Protection Reports **Delegation** Settings Support

Help Desk Delegation

- Help Desk Technicians
- Help Desk Roles
- Help Desk Audit Reports
 - Audit Report
 - Admin Audit Report
- Configuration
- Help Desk Reports

Help Desk Technicians

Create help desk technicians and delegate the desired tasks/roles to them. [Learn more...](#)

Q Manage

☐	Action	Workflow Role	Name	↑
☐	  	Requester, Reviewer, Approver, Executor	 SecureAI Admin	
☐	  	Approver, Executor	 Governance Technician	
☐	  	Reviewer	 Gateway Admin	
☐	  	Reviewer, Approver, Executor	 DemoAccount1	
☐	  		 DemoAccount2	
☐	  	Reviewer, Approver, Executor	 DemoAccount3	
☐	  	Reviewer, Approver, Executor	 DemoAccount4	
☐	  	Reviewer, Approver, Executor	 DemoUser	

Add New Technician

Authentication Method *

Select Tenant *

Select Microsoft 365 Users *

Select Roles *

Delegate without over-granting

Dashboard

Governance

Access Protection

Reports

Delegation

Settings

Support

Application Settings

Help Desk Technician

Help Desk Roles

Audit Report

Admin Audit Report

Help Desk Roles

*Role Name

Description

☒ Governance

☒ Access Protection

☐ Reports

☒ Administration

☒ AI Inventory

☒ AI Agents

☒ MCP Servers

☒ Models

☒ Unmanaged Agents

☒ High Privileged Agents

☒ Access Configuration

☒ Ownership Policy

☒ Entitlement

☒ Access Templates

☒ Risk Management

☒ AI Risk Assessment

☒ Access Certification

☒ Access Review Request

☒ Access Certification Campaign

☒ Governance Assigning Rule

Multi-platform reporting

Pre-built report packs per platform: Microsoft Agents, Gemini Enterprise Agent Platform, and Zia Agents—agent inventories with recent, unmanaged, and high-privilege cuts, plus permission and capability reports for each. Microsoft adds a blocked-agents view.

The screenshot displays the ManageEngine Reports interface. The top navigation bar includes links for Dashboard, Governance, Access Protection, Reports (active), Delegation, Settings, and Support. A search bar and a dropdown menu for platform selection (Microsoft Agent, Gemini Enterprise Agent Platform, More) are located below the navigation bar. The left sidebar lists the platforms: Microsoft Agent (selected), Gemini Enterprise Agent Platform, and Zoho Zia Agents. The main content area is titled "Microsoft Agent" and is divided into three columns of report packs:

- Agents**
 - All Agents
 - Copilot Studio Agents
 - AI Foundry Agents
 - Recently Created Agents
 - Recently Modified Agents
- Agent Governance Report**
 - Blocked Agents
 - Unmanaged Agents
 - High Privileged Agents
- Agent Capability**
 - All Tools (Copilot Studio)
 - All Tools
 - All Knowledge Bases
 - All Guardrails

SecureAI at a glance

Discover	Every agent, model, and MCP server across Microsoft, Google, and Zia
Govern	Ownership, entitlements, access templates, and recurring certification
Protect	An MCP gateway that inspects and controls every tool call at runtime
Assess	Misconfiguration risk scoring for Anthropic Claude, GitHub Copilot, and Gemini
Operate	Scoped delegation and per-platform reporting

One platform, instead of stitching together discovery, IAM, runtime control, and compliance separately.



Thank you