Introduction to Agent Workforce

This chapter introduces Agent Workforce, an AI-driven operational framework in Security Cloud Control that helps you troubleshoot network issues, analyze operational behavior, and manage firewall policies using natural language interactions.

Agent Workforce uses specialized agents to assist with operational workflows such as VPN troubleshooting, traffic remediation, and policy analysis. You can interact with agents conversationally to investigate issues, review recommendations, and perform guided operational tasks.

- About Agent Workforce , on page 1
- Troubleshoot site-to-site VPN issues with VPN agent, on page 3
- Remediate traffic congestion with Elephant flow agent , on page 6
- Optimize firewall policies with Policy Copilot , on page 7

About Agent Workforce

Agent Workforce is an AI-driven operational framework that enables you to interact with specialized agents using natural language to troubleshoot issues, analyze network behavior, and manage firewall policies.

Instead of relying solely on manual investigation, command-line analysis, or complex configuration workflows, Agent Workforce allows you to describe operational problems or intended actions conversationally. Based on the user query, the appropriate agent analyzes telemetry, configurations, policies, and operational data to provide insights, recommendations, diagnostics, or policy actions. Each agent is optimized for a specific operational area and performs focused analysis based on its specialization.

Available agents in Agent Workforce

Agent Workforce includes specialized agents designed to assist with different operational and security workflows. Each agent focuses on a specific operational domain and uses relevant telemetry, protocol awareness, policy analysis, and contextual operational data to provide diagnostics, recommendations, and actionable outcomes.

When you enter a query, Agent Workforce interprets the intent and automatically routes the conversation to the most appropriate agent. You can describe operational issues or tasks in natural language without determining which agent should handle the request.

Note Currently, Agent Workforce capabilities are available only for deployments in the AMER and EU regions.

Depending on the workflow, Agent Workforce may also present contextual quick actions to help you initiate common investigative or remediation tasks more efficiently. Currently, Agent Workforce includes the following agents:

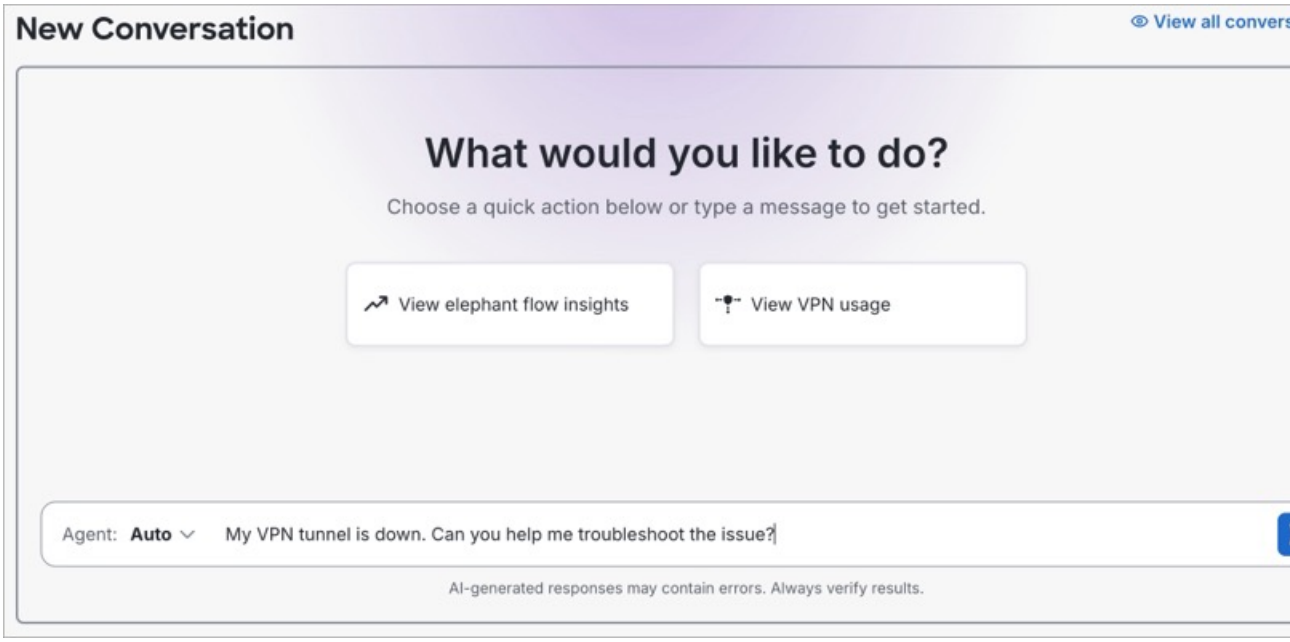

- **VPN agent** helps troubleshoot site-to-site VPN connectivity, routing, tunnel negotiation, and secure traffic flow issues.
- **Elephant flow agent** helps analyze and remediate high-bandwidth traffic flows that may impact firewall performance and operational stability.
- **Policy Copilot** helps create, analyze, modify, and optimize firewall policies while providing policy insights, validation, and optimization guidance.

Manage Conversations in Agent Workforce

Use the **Conversations** page to create new conversations, monitor active investigations, and interact with specialized agents for troubleshooting, remediation, and policy management tasks.

Procedure

	Command or Action	Purpose
Step 1	In the left pane, click **Insights & Reports** > **Agent Workforce** > **Conversations.**	

	Command or Action	Purpose
Step 2	Review the list of available conversation threads.	Each conversation displays information such as the thread title, assigned agent, workflow status, creator, and creation date.
Step 3	To locate a specific conversation, use the **Search** field to search by keywords.	
Step 4	To narrow the conversation list, click **Filters** and filter conversations using one or more criteria.	
Step 5	To create a new conversation, click **New Conversation**.	
Step 6	Enter your query in natural language.	Agent Workforce automatically selects the most appropriate agent based on the operational intent of the request.
Step 7	Continue interacting with the assigned agent if additional investigation, troubleshooting, or remediation actions are required.	

Limitations and considerations

- Currently, Agent Workforce capabilities are available only for deployments in the AMER and EU regions.
- Ensure that Cloud-Delivered Firewall Management Center is provisioned in your organization.
- AI-generated insights, recommendations, and remediation actions must always be reviewed and validated before deployment.
- The accuracy of generated responses depends on the availability, completeness, and correctness of telemetry, policies, device configurations, logs, objects, and operational context within the environment.
- Missing, incomplete, or unresolved configuration entities may prevent certain workflows, recommendations, remediation actions, or deployment operations from completing successfully.
- Agent responses and operational analysis may vary depending on the enabled platform capabilities, configured integrations, available data sources, and device reachability.
- Some workflows may require manual validation, administrative approval, or additional configuration updates before operational changes can be applied to managed devices or policies.
- AI-generated operational guidance is intended to assist administrators and must not replace standard security review, compliance validation, or change-management procedures.

Troubleshoot site-to-site VPN issues with VPN agent

The VPN agent is an intelligent troubleshooting agent that helps you diagnose and resolve site-to-site VPN connectivity and access issues. The agent analyzes operational telemetry, VPN configurations, routing behavior, packet flow, and access control policies to identify issues affecting secure connectivity between sites. It

provides protocol-aware diagnostics across technologies such as BGP, EIGRP, OSPF, IPsec, and IKE to accelerate root-cause analysis and reduce manual troubleshooting effort.

Benefits

- Investigate VPN tunnel failures and connectivity disruptions.
- Analyze traffic flow between remote networks.
- Detect routing inconsistencies affecting VPN communication.
- Identify access control or policy-related packet drops.
- Perform packet-level troubleshooting and operational analysis.

How VPN agent works

When a query is submitted, the VPN agent correlates operational data across routing protocols, tunnel negotiation states, traffic flow behavior, and policy enforcement to identify the most likely root cause and provide actionable insights. Depending on the issue being investigated, the VPN agent can:

- Identify affected site-to-site VPN tunnels.
- Analyze IKE Phase 1 and Phase 2 negotiation states.
- Review tunnel establishment and peer connectivity status.
- Investigate routing behavior across the VPN tunnel.
- Evaluate traffic flow behavior between local and remote networks.
- Correlate logs, telemetry, and operational events to identify likely root causes.
- Generate operational findings, root-cause analysis, and remediation recommendations.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Agent Workforce** > **Conversations**.

Step 2 Click **New Conversation**.

Step 3 Enter a query describing the VPN connectivity issue or traffic-flow problem.

Agent Workforce automatically assigns the request to the VPN agent based on the operational intent.

Step 4 Review the generated diagnostics, tunnel analysis, routing behavior, traffic flow analysis, and remediation recommendations.

Step 5 Continue refining the investigation by providing additional operational details such as affected sites, peer devices, tunnel names, or observed behavior.

Step 6 Review the recommended remediation actions and validate the operational impact before applying configuration changes.

After the analysis is completed, the VPN agent summarizes operational findings, identified root cause, impact assessment, additional observations, recommended remediation steps and next actions.

Best practices

To improve troubleshooting accuracy and reduce analysis time:

- Include source and destination networks whenever possible.
- Specify the affected VPN tunnel, site, or peer device if known.
- Clearly describe observed behavior such as packet drops, intermittent connectivity, or failed tunnel establishment.
- Refine prompts with additional operational context if further analysis is needed.
- Validate recommended remediation steps before deployment.

Example prompts and expected outcomes

User intent	Example prompt	Expected outcome
Troubleshoot connectivity between remote sites	Troubleshoot why the connectivity is down between [remote site 1] and [remote site 2].	VPN agent identifies the relevant Site-to-Site VPN tunnel associated with the affected networks and performs guided troubleshooting analysis.
Troubleshoot a VPN tunnel outage	My VPN tunnel is down. Help me troubleshoot the issue.	VPN agent identifies available VPN tunnels, allows tunnel selection if needed, and analyzes tunnel state, negotiation behavior, routing, and connectivity conditions.
Investigate VPN traffic flow issues	Investigate why VPN traffic is not flowing across the [tunnel name].	VPN agent analyzes traffic forwarding behavior, routing conditions, and policy enforcement affecting traffic flow across the VPN tunnel.
Review configured VPN tunnels	How many VPN tunnels are configured and what state are they in?	VPN agent displays configured site-to-site VPN tunnels along with operational state, tunnel health, and connectivity status details.
Visualize VPN routing paths	Visualize the routing between various devices connected by my VPN tunnels.	VPN agent analyzes VPN topology, connected networks, and routing relationships, and provides communication paths across VPN tunnels.

Remediate traffic congestion with Elephant flow agent

The Elephant flow agent is an intelligent remediation agent that helps you analyze and remediate previously detected high-bandwidth traffic flows that impact network performance and critical applications. Large traffic flows, commonly referred to as elephant flows, can consume excessive bandwidth, increase firewall processing overhead, and degrade application performance across the network. Unlike traditional troubleshooting workflows that require manual investigation across multiple dashboards and telemetry sources, the Elephant flow agent provides a guided remediation experience directly within the conversation workflow.

Benefits

- Analyze previously detected elephant flows affecting network performance.
- Investigate abnormal traffic spikes and sustained utilization.
- Identify applications or traffic flows contributing to congestion.
- Recommend remediation actions to reduce processing overhead and congestion.
- Stage policy updates associated with remediation workflows.

How Elephant flow remediation works

Procedure

Step 1 In the left pane, click **Insights & Reports** >**Summary**.

Step 2 On the AIOps **Summary** page, select the **Traffic & Capacity** insights category.

Step 3 In the **High traffic caused by elephant flow** section, select the affected device.

Step 4 Review the detected elephant flow insight details, affected resources, applications, and traffic impact information.

Step 5 Click **Review proposed remediation** to open the remediation workflow in **Agent Workforce** > **Conversations**.

a. Alternatively, navigate to **Insights & Reports > Agent Workforce > Conversations**.

b. Click **New Conversation** and select **View elephant flow insights** to start the remediation workflow.

Note
Manual prompt entry is not required to start the workflow.

Step 6 Review the proposed remediation actions, operational warnings, shared policy impact, and success criteria.

Step 7 Approve or cancel the remediation workflow as required.

Step 8 Deploy the generated policy updates to the affected devices.

Best practices

To improve remediation accuracy and operational visibility:

- Review affected applications before approving remediation actions.
- Verify the operational impact of policy updates before deployment.
- Review shared policy impact warnings before approving remediation actions.
- Monitor application and network behavior after remediation is completed.

Optimize firewall policies with Policy Copilot

Policy Copilot is an AI-driven assistant that helps simplify firewall policy creation, analysis, optimization, and operational validation using natural language interactions.

Instead of manually navigating policy objects, rule placement, logging options, inspection policies, and deployment workflows, you can describe the intended access behavior conversationally. Policy Copilot interprets the request, analyzes the available policy context, validates referenced entities, and generates recommended policy configurations.

Policy Copilot also provides operational guidance, conflict analysis, and deployment readiness checks before policies are approved or deployed.

Benefits

- Identify and validate referenced networks, hosts, applications, services, zones, and policy objects.
- Recommend suitable access control policies for rule placement.
- Support iterative refinement of generated policies through conversational interactions.
- Validate generated rules against existing policies to identify overlaps, conflicts, shadowed rules, duplicates, and redundant entries before deployment.
- Generate multiple options with different security postures and monitoring configurations.
- Provide policy explanations, operational impact details, and security considerations.
- Recommend intrusion, inspection, and logging configurations when applicable.
- Detect missing or unresolved policy objects required for deployment.
- Organize generated rules using policy affinity analysis based on zones, networks, services, protocols, applications, and action types.
- Capture business justification during rule creation, update, and deletion workflows to improve policy intent visibility and change traceability.
- Identify policy drift when rule modifications deviate from the original rule purpose, approved access behavior, or business justification.

How Policy Copilot works

When a query is submitted, Policy Copilot analyzes the intent, evaluates the existing policy environment, and generates one or more policy recommendations based on the requested access behavior.

Note

- Policy Copilot displays up to 20 results at a time for policy, rule, and object-related queries.
- To view additional results, use prompts such as “Show next 20” or “Show next 20 network objects”.
- Policy Copilot remembers the current policy context and active filters during follow-up interactions, so you do not need to repeat them in subsequent queries.
- Policy Copilot can capture business justification for rules created through **Conversations** only.
- Supported object types include:
  - **Network Object**: Host (single IP), Range (IP range), Network (CIDR), and FQDN
  - **Port Object**: TCP and UDP port definitions
  - **Network Group Object**: Collection of network objects and inline IP/CIDR literals
  - **Port Group Object**: Collection of named port objects

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Agent Workforce** > **Conversations**.

Step 2 Click **New Conversation**.

Step 3 Enter a query describing the required policy behavior.

Policy Copilot automatically analyzes the request and extracts relevant rule requirements.

Step 4 Review the available policy recommendations and generated rule options.

Step 5 Provide refinements or additional requirements to modify the generated recommendations if needed.

Step 6 Review any warnings related to missing objects, deployment readiness, overlaps, or policy conflicts.

Step 7 Approve the preferred option to continue with deployment preparation.

Object creation limitations and considerations

- Policy Copilot supports creation of one object per request. Bulk object creation is not supported.
- URL objects, SGT objects, Application objects, GeoLocation objects, and unsupported protocol-specific port objects are not currently supported.
- Network groups support inline IP addresses and CIDR entries during creation workflows.
- Port group members must reference existing named port objects. Inline port values are not supported for port groups.

- Nested group references are supported only for existing groups. Creating a new group inline within another group is not supported.

Best practices

To improve troubleshooting accuracy and reduce analysis time:

- Provide detailed context such as applications, networks, services, zones, access requirements, or intended policy behavior whenever possible to improve recommendation accuracy and policy relevance.
- Use follow-up questions to refine policy analysis and investigations. For example, review a policy, examine its rules, analyze specific details, and explore related categories or objects within the same conversation workflow.
- Review all generated policy recommendations before approval and deployment.
- Validate logging and inspection settings to ensure the required visibility and security coverage.
- Resolve missing objects and policy conflicts before deployment.
- Use conversational refinements to improve accuracy.
- Verify generated policies in the Firewall Management Center before deploying changes to production environments.

Example prompts and expected outcomes

User intent	Example prompts	Expected outcome
View policies	• List all policies • List all policies with rule counts • List all policies and show which devices they are assigned to	Displays available access control policies and related information such as assigned devices and rule counts.
Review policy details	• Describe the policy [policy name] • Summarize the [policy name] policy • Give me an overview of what [policy name] does	Provides a high-level summary of the selected policy, including its purpose and rule composition.

User intent	Example prompts	Expected outcome
Search and filter rules	• List all “Allow” rules in [policy name] • Find rules using port 22 in [policy name] • Show rules with source zone External in [policy name]	Retrieves rules matching the specified filters, criteria, or policy scope.
Create rules	• Create a temporary “Allow” rule in [policy name] to allow [source object] to access [destination subnet] over HTTP and SSH • Create an “Allow” rule in [policy name] for HTTPS traffic and attach intrusion and file policies with a business justification	Generates policy recommendations, validates referenced objects, recommends rule placement, and captures business justification for the requested policy change.
Modify rules	• Update rule [rule name] to remove HTTPS access • Modify rule [rule name] action to “Deny”	Updates the selected rule, validates policy impact, identifies conflicts or policy drift, and recommends additional verification if required.
Delete rules	• Delete rule [rule name]	Identifies the rule, validates dependencies and policy impact, and analyzes operational considerations before deletion.
Detect expiring rules	• Find expiring rules in [policy name] • Describe these expiring rules • Find rules expiring in the next 90 days in [policy name]	Identifies rules associated with expiring schedules and provides lifecycle management recommendations.
Review policy activity and intent	• Why does rule [rule name] exist? • What is the business purpose of rule [rule name]?	Retrieves rule history, business justification, and policy intent information.

User intent	Example prompts	Expected outcome
Work with policy categories	• List all categories in [policy name] • Show rules in the DMZ category of [policy name] • List rules in the outbound category of [policy name]	Displays available policy categories and retrieves rules associated with the selected category.
Cross-policy discovery	• Find policies that reference object [object_name] • Find policies containing object [object_name] • Find rules containing keyword “AI” across all policies • Find rules for IP address 192.168.10.1	Searches across multiple policies and identifies matching rules, objects, references, or IP-based relationships.
Review memory and activity	• What activities did I perform? • Search my activities for [policy_name] • Did I recently work on any rules for the DMZ zone?	Retrieves conversationally captured rule intent, business justification, recent activity, and previously modified policy information.
Search objects	• List all network objects • Find port object for port 443 • Find object "example_object"	Searches supported object types and returns matching configuration objects.
Create network objects	• Create a host object for `10.10.10.10` named `Branch_Server` • Create a network object for `10.20.0.0/16` • Create an FQDN object for `example.com`	Creates the requested supported network object and validates naming and object constraints.
Create port objects	• Create a TCP port object for port 443 • Create a UDP port object for port 53	Creates supported TCP or UDP port objects for policy usage.

User intent	Example prompts	Expected outcome
Create network groups	• Create a network group containing `10.10.10.0/24` and `Branch_Server` Create a network group for branch office networks	Creates a network group using supported object references or inline network literals.
Create port groups	• Create a port group containing `HTTPS` and `SSH` • Create a port group for web applications	Creates a port group using existing named port objects.