

INTELLIGENCE BRIEFING

Security Command Center

TLP:CLEAR

2026-05-09 18:47 UTC

Solo.io Integrates NemoClaw Governance Framework into kagent for Secure AI Agent Deployment on Kubernetes

GOVERNANCE | LOW

SCC Item ID	SCC-GOV-2026-0032
Type	Governance
Severity	LOW
Affected Products	Solo.io kagent runtime, Kubernetes-hosted AI agents, NemoClaw framework (CNCF ecosystem)
Published	2026-05-08
Discovery Source	Gemini

Executive Summary

Solo.io has integrated the NemoClaw governance framework into its open-source kagent runtime, adding least-privilege access controls and context boundary enforcement for AI agents running on Kubernetes. Organizations deploying AI agents in production Kubernetes environments gain a structured mechanism to prevent agents from accessing data or resources beyond their defined scope. This is a proactive governance addition to address a common architectural gap in AI agent deployments; business risk is low for informed adopters and moderate for organizations running AI agents without equivalent controls.

Technical Analysis

Solo.io's kagent is a Kubernetes-native runtime for executing AI agents in production environments. The NemoClaw integration, a community-driven governance framework in the CNCF ecosystem, enforces least-privilege data access (addressing CWE-732: Incorrect Permission Assignment for Critical Resource) and context boundary controls (addressing CWE-285: Improper Authorization) at the agent execution layer. Without these controls, AI agents can exhibit scope creep, accessing data stores or compute resources beyond their intended operational boundary, mapping to MITRE ATT&CK T1078 (Valid Accounts, where an agent's legitimate credentials enable unintended access) and T1530 (Data from Cloud Storage, where over-permissioned agents can read cloud-resident data). No CVE is assigned. No patch is required; this is an opt-in governance capability. Organizations using kagent should evaluate current agent permission scopes against NemoClaw policy definitions. Organizations using other Kubernetes-based AI agent runtimes should assess equivalent control coverage independently. Source quality is T3 (vendor and trade press); no primary-source CNCF or NIST documentation for NemoClaw was available at time of writing.

Action Checklist

1. Step 1: Assess Current State, Inventory all AI agents running in your Kubernetes clusters, document their current service account permissions, and map which data stores and APIs each agent can reach today.
2. Step 2: Audit Existing Controls, Review Kubernetes RBAC bindings for AI agent service accounts; flag any service account with cluster-wide read permissions or access to secrets, ConfigMaps, or external storage beyond its documented operational need.
3. Step 3: Implement Governance Framework, If using kagent, evaluate the NemoClaw governance integration and define explicit context boundary policies per agent; if using a different runtime, apply equivalent least-privilege RBAC scoping and network policy restrictions.
4. Step 4: Validate and Monitor, Validate agent functionality post-scoping with integration tests; monitor agent logs for authorization errors that indicate over-restriction, and adjust policies with minimum necessary permissions.
5. Step 5: Formalize AI Agent Governance, Document approved data access scopes per agent role, integrate agent permission reviews into your existing access recertification cycles, and align with NIST AI RMF Govern and Manage functions.

IR / Forensic Enrichment

Triage Priority	DEFERRED
Escalation Criteria	Escalate to urgent if discovery in Step 2 reveals that any AI agent service account has current or historical access to Kubernetes secrets containing credentials, API keys, or PII-bearing ConfigMaps — at that point the exposure becomes a potential data access incident requiring breach notification assessment under applicable regulatory obligations.
Recovery Notes	Post-scoping, monitor Kubernetes API server audit logs and agent pod logs for a minimum of 72 hours to confirm NemoClaw context boundary policies or custom RBAC Roles are enforcing correctly without breaking documented agent workflows. Verify that no agent pod has been restarted with a cached token predating the RBAC change by checking pod creation timestamps against the remediation timestamp using <code>kubectl get pods -A -o wide`</code> . Treat any recurrence of cluster-wide secret access by an agent service account as a policy bypass event requiring immediate re-investigation.

Forensic Artifacts	Kubernetes API server audit log entries (filtered by agent service account subjects) showing `get`, `list`, or `watch` verbs against `secrets`, `configmaps`, or `persistentvolumeclaims` — primary evidence of pre-remediation access scope for kagent or other AI agent runtimes RBAC YAML exports (`kubectl get clusterrolebindings,rolebindings -A -o yaml`) captured before and after NemoClaw policy or RBAC scoping changes — documents the permission delta and serves as the access recertification baseline Agent pod stdout/stderr logs from the kagent runtime, specifically `Forbidden` and `RBAC: access denied` entries generated during recovery validation — evidence of correct context boundary enforcement post-remediation NemoClaw context boundary policy objects or custom Kubernetes Role/RoleBinding manifests as applied in the remediation step — configuration artifacts proving least-privilege scoping was implemented per agent identity `kubectl auth can-i --list --as=system:serviceaccount::` output snapshots per agent service account, timestamped pre- and post-remediation — point-in-time permission surface evidence for audit and recertification purposes
---------------------------	--

Per-Action IR Details

Step 1: Inventory — identify all AI agents running in your Kubernetes clusters, document their current service account permissions, and map which data stores and APIs each agent can reach today.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing asset visibility and access baselines before incidents occur

Controls: NIST IR-4 (Incident Handling) — preparation phase requires current asset and permission baselines, NIST SI-4 (System Monitoring) — monitoring scope depends on knowing what agents exist and what they access, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — extend asset inventory to include kagent and other AI agent workloads as distinct asset types, CIS 5.1 (Establish and Maintain an Inventory of Accounts) — service accounts bound to AI agent pods in Kubernetes are in-scope accounts requiring documentation

Compensating: Run `kubectl get serviceaccounts -A` to list all service accounts across namespaces, then `kubectl get rolebindings,clusterrolebindings -A -o json | jq '.items[] | select(.subjects[]?.name | test("agent|kagent|ai"))'` to filter bindings associated with AI agent identities. Cross-reference with `kubectl auth can-i --list --as=system:serviceaccount::` per discovered account to enumerate effective permissions. Export results to a CSV for baseline documentation. No SIEM required.

Evidence: Before scoping changes, capture the current permission baseline: export `kubectl get clusterrolebindings -o yaml` and `kubectl get rolebindings -A -o yaml` to preserve the pre-change RBAC state. Also snapshot `kubectl get pods -A -o json` to record which pod names are running agent workloads and which service accounts they are bound to — this establishes the before-state for any future access creep investigation.

Step 2: Detection — audit Kubernetes RBAC bindings for AI agent service accounts; flag any service account with cluster-wide read permissions or access to secrets, ConfigMaps, or external storage beyond its documented operational need.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection & Analysis: identifying indicators of misconfiguration and excessive privilege that represent realized risk

Controls: NIST AU-6 (Audit Record Review, Analysis, and Reporting) — RBAC audit logs and Kubernetes API server audit logs must be reviewed for anomalous access patterns by agent service accounts, NIST SI-4 (System Monitoring) — monitor Kubernetes API server audit log for agent service accounts accessing secrets, ConfigMaps, or cross-namespace resources outside documented scope, NIST IR-5 (Incident Monitoring) — track and document flagged over-privileged bindings as candidate incidents pending context boundary policy definition, CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts) — cluster-wide read on secrets by an AI agent service account is an equivalent privilege violation in the Kubernetes context, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — overly permissive RBAC for AI agents is a configuration risk finding requiring tracked remediation

Compensating: Enable Kubernetes API server audit logging if not already active (set `--audit-log-path` and `--audit-policy-file` in kube-apiserver manifest). Then query the audit log with: `grep -E "user".*agent|"serviceAccount" /var/log/kubernetes/audit.log | jq 'select(.objectRef.resource=="secrets" or .objectRef.resource=="configmaps")'` to surface agent access to sensitive resources. Use `kubectl-who-can` (open-source, <https://github.com/aquasecurity/kubectl-who-can> — recommend human validation of URL) to enumerate which service accounts can perform `get secrets` cluster-wide. Flag any result not matching documented operational need.

Evidence: Capture Kubernetes API server audit logs (`/var/log/kubernetes/audit.log` or cloud provider equivalent — e.g., GKE Audit Logs, EKS CloudTrail `kubernetes.io` events) filtered for agent service account subjects performing `get`, `list`, or `watch` on `secrets`, `configmaps`, or `persistentvolumeclaims`. Also preserve `kubectl describe clusterrolebinding` output for any binding granting `cluster-admin` or wildcard resource access to an agent service account — these are the primary evidence of scope violation before NemoClaw context boundaries are applied.

Step 3: Eradication — if using kagent, evaluate the NemoClaw governance integration and define explicit context boundary policies per agent; if using a different runtime, apply equivalent least-privilege RBAC scoping and network policy restrictions.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication: removing the conditions that allowed the risk to exist, in this case over-broad RBAC and absent context boundary enforcement

Controls: NIST SI-2 (Flaw Remediation) — applying NemoClaw governance policies or equivalent RBAC scoping closes the architectural gap identified in the detection phase, NIST CM-6 (Configuration Settings) — context boundary policies in NemoClaw and Kubernetes NetworkPolicy objects are configuration settings requiring documented baselines, CIS 4.4 (Implement and Manage a Firewall on Servers) — Kubernetes NetworkPolicy is the equivalent server-level firewall control; apply egress restrictions on agent pods to named services only, CIS 7.2 (Establish and Maintain a Remediation Process) — RBAC scoping changes for each agent service account should be tracked in a risk-based remediation log with before/after states

Compensating: For kagent environments, follow Solo.io's NemoClaw integration documentation to define per-agent `ContextBoundary` policy objects specifying allowed namespaces, secrets, and API endpoints. For non-kagent runtimes, create a dedicated Role per agent namespace with explicit resource allowlists: `kubectl create role agent--role --verb=get,list --resource=configmaps --namespace=` and bind it with a RoleBinding (not ClusterRoleBinding). Apply Kubernetes NetworkPolicy to restrict egress from agent pods to only the specific service IPs or DNS names they require — use `kubectl apply -f` with a deny-all-egress default plus explicit allow rules per agent.

Evidence: Before applying NemoClaw policies or RBAC changes, export the full current RBAC state: `kubectl get clusterrolebindings,rolebindings -A -o yaml > rbac-before-$(date +%F).yaml`. This preserves the over-privileged baseline as a forensic reference. Also document which agent pods had access to which secrets by name (`kubectl get secret -A --field-selector metadata.namespace=`) — this establishes what data was within scope of the over-broad permissions prior to remediation.

Step 4: Recovery — validate agent functionality post-scoping with integration tests; monitor agent logs for authorization errors that indicate over-restriction, and adjust policies with minimum necessary permissions.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery: restoring systems to verified operational state after eradication, confirming controls do not break legitimate function

Controls: NIST IR-4 (Incident Handling) — recovery phase requires verification that remediation did not introduce operational degradation, NIST AU-2 (Event Logging) — agent pod logs and Kubernetes API server audit logs must be actively monitored post-scoping for denied access events indicating policy misconfiguration, NIST SI-6 (Security and Privacy Function Verification) — verify that NemoClaw context boundary policies or custom RBAC Roles are enforced as intended by running agents through their documented workflows and confirming no unexpected access is permitted, CIS 8.2 (Collect Audit Logs) — ensure agent pod stdout/stderr logs and Kubernetes audit logs are being collected and retained during the post-change validation window

Compensating: Monitor agent pod logs in real time during validation: ``kubectl logs -f -n `` filtering for HTTP 403 responses or ``RBAC: access denied`` messages. Query the Kubernetes audit log for ``DENY`` verdicts against the agent service account: ``grep -E 'Forbidden|403' /var/log/kubernetes/audit.log | jq 'select(.user.username | test("system:serviceaccount:"))'``. If authorization errors appear for legitimate operations, use ``kubectl auth can-i --as=system:serviceaccount:`` to diagnose and surgically expand the Role — do not revert to ClusterRoleBinding.

Evidence: Capture agent pod logs and Kubernetes API server audit log entries for the 24–48 hours immediately following policy application, specifically filtering for the agent service account names. Preserve any ``Forbidden`` or ``401 Unauthorized`` events as evidence of correct enforcement. Also run ``kubectl auth can-i --list --as=system:serviceaccount:`` post-remediation and save the output to confirm the reduced permission surface — this serves as the post-remediation baseline for future access recertification audits.

Step 5: Post-Incident — treat this as a prompt to formalize AI agent governance policy: document approved data access scopes per agent role, integrate agent permission reviews into your existing access recertification cycles, and align with NIST AI RMF Govern and Manage functions.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: lessons learned, policy updates, and capability improvements to prevent recurrence

Controls: NIST IR-8 (Incident Response Plan) — update the IR plan to include AI agent governance as an explicit scope item, referencing NemoClaw or equivalent runtime controls as standard tooling, NIST IR-2 (Incident Response Training) — train relevant personnel on AI agent RBAC review procedures and NemoClaw context boundary policy authoring as part of the next IR training cycle, NIST RA-3 (Risk Assessment) — document residual risk for any agent service account that required exceptions to strict least-privilege during recovery; schedule reassessment at next policy review, CIS 6.1 (Establish an Access Granting Process) — extend the access granting process to include AI agent service account provisioning with explicit scope documentation as a required artifact, CIS 6.2 (Establish an Access Revoking Process) — integrate AI agent service account review into the existing access recertification cycle, with a trigger for review on any new agent deployment or kagent version upgrade

Compensating: Create a Markdown or YAML policy document per agent role defining: allowed namespaces, allowed secrets by name, allowed ConfigMaps, allowed external endpoints, and maximum RBAC verbs. Store this in version control alongside the agent deployment manifests so permission scope is code-reviewed on every change. Schedule a quarterly ``kubectl auth can-i --list`` audit per agent service account using a cron job or CI pipeline step — compare output against the documented policy baseline and alert on any divergence. No SIEM required; a shell script and git diff are sufficient.

Evidence: As a post-incident record, archive the before-and-after RBAC YAML exports from Steps 1 and 4, the NemoClaw context boundary policy objects or custom Role manifests applied in Step 3, and the post-remediation ``kubectl auth can-i --list`` outputs per agent service account. These collectively document the permission reduction achieved and serve as the audit trail for the first access recertification cycle. Also capture the Kubernetes API server audit log segment covering the full remediation window for retention per NIST AU-11 (Audit Record Retention).

Detection Guidance

Review Kubernetes audit logs for service account activity associated with AI agent workloads. Query for events where agent service accounts access Secrets, ConfigMaps, PersistentVolumes, or external storage endpoints (S3-compatible APIs, cloud object storage) outside their expected namespace. In cloud environments, enable data access logging on storage buckets and flag reads initiated by Kubernetes workload identities that do not correspond to defined agent functions. There are no IOCs for this item; detection focus is misconfiguration and over-permissioning, not active exploitation.

Framework Mappings

MITRE-ATTACK

- **T1078** — Valid Accounts
- **T1530** — Data from Cloud Storage

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-3** — Access Enforcement

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control

CIS-V8

- **3.3** — Configure Data Access Control Lists

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1078	Valid Accounts	Defense-Evasion
T1530	Data from Cloud Storage	Collection

Sources

Source	URL	Tier
gemini	https://cloudnativenow.com/topics/kubernetes/solo-io-extends-kagent...	T3
Kagent: Kubernetes-Native Agent Runtime - Solo.io	https://www.solo.io/products/kagent	T3
Securing AI Agents & Tools in Kubernetes Solo.io - YouTube	https://www.youtube.com/watch?v=9IEsKNCd8S4	T3
Solo.io announces kagent enterprise - ITOps Times	https://itopstimes.com/agentic-ai/solo-io-announces-kagent-enterprise/	T3

Source	URL	Tier
Inside AI Agents in Kubernetes: Live with the Authors - Webinar	https://www.solo.io/resources/webinar/inside-ai-agents-in-kubernetete...	T3

DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness. Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-05-09 18:47 UTC by TJS Security Command Center