Illumio® Core for Kubernetes

Version 5.1.7
Compatible PCE Versions: 23.5.10 and higher

Release Notes

07/01/2024

19750-100-5.1.7

Contents

Welcome

These release notes describe the resolved issues, known issues, and related information for the 5.1.*x* releases of Illumio Core for Kubernetes, formerly known as Illumio Containerized VEN, or C-VEN. Illumio Core for Kubernetes also includes the related required component, Kubelink. Because of this heritage, many references to this product as "C-VEN" are still used throughout the documentation.

Document Last Revised: July 2024
Document ID: 19750-100-5.1.7

Product Version

Compatible PCE Versions: 23.5.10 and most later releases
Current Illumio Core for Kubernetes Version: 5.1.7, which includes:

- C-VEN version: 23.3.0
- Kubelink version: 5.1.7
- Helm Chart version: 5.1.7

Before deploying any Illumio Core for Kubernetes 5.1.x version, confirm your PCE version supports it. For example, currently Illumio Core for Kubernetes versions 5.1.0 and 5.1.2 are supported **only** with PCE versions 23.5.10 (for On Premises customers) or 24.1.x (for SaaS customers), but NOT on PCE versions 23.5.1 or 23.6.0, or any lower versions. For complete compatibility details, see the **Kubernetes Operator OS Support and Dependencies** page on the Illumio Support Portal.

Release Types and Numbering

Illumio release numbering uses the following format: “a.b.c-d”

- “a.b”: Standard or LTS release number, for example, “21.2”
- “.c”: Maintenance release number, for example, “.1”
- “-d”: Optional descriptor for pre-release versions, for example, “preview2”

Limitations

- **NodePort**
  The following limitations exist regarding NodePort policy enforcement and flows:
  - Only NodePort Services with `externalTrafficPolicy` set to "`cluster`" are supported. (This is the default and most frequently used value for this setting.)
  - When writing rules to allow traffic to flow from external (to the cluster) entities and NodePort Service, the source side of the rule must contain all nodes in the cluster. For example, given the following setup:
    - Worker nodes in the cluster are labeled as Role: Worker Node
    - Clients accessing the Service running in the Kubernetes cluster are labeled Role: Client
    - The NodePort Service is labeled Role: Ingress

    Normally, the rule would be written as Role: Client -> Role: Ingress. However, for this release the rule must also include all nodes in the cluster to work correctly: Role: Client + Role: Worker Node -> Role: Ingress.
- **Flat Network support in CLAS mode**
  Using EKS or AKS in a flat network topology, such as EKS with AWS VPC CNI or AKS with Azure CNI, is not supported in CLAS-enabled clusters.

Updates for Core for Kubernetes 5.1.7

Kubelink

Resolved Issues

- **Kubelink: policy service blocked when agent disconnects while receiving policy message** (E-117099)
  In some situations, policies stopped being sent due to a policy channel lock after C-VEN disconnected while receiving a policy update.
- **Kubelink: policy service blocked if one agent is not reading policy message** (E-116967)
  In some situations, policies stopped being sent after a C-VEN became unresponsive.
- **Kubelink can't save sets because of message size limit** (E-116825)
  Policy updates were being interrupted when large policy sets were being sent. The message size has been increased to permit larger policy transmissions.
- **Kubelink: workload events processing is slowed down by policy updates** (E-116706)
  The processing of workload events from Kubernetes sometimes became slow when handling

thousands of Kubernetes Workloads, or the policy PCE requests were taking too long, or if there was no previous policy version in storage.

- **Kubelink sends wrong workload href in policy ACK request** (E-116640)
  In some CLAS-enabled clusters that host large numbers of workloads, the Kubernetes Workloads page showed an old policy apply date. Kubelink incorrectly sent a policy ACK for some Kubernetes Workloads with the host workload URI. The PCE responded with a 406 error, and a "no policy" ACK was stored.

Updates for Core for Kubernetes 5.1.3

Kubelink

Resolved Issues

- **Kubelink can't save policy to storage** (E-116539)
  Kubelink could not store cluster policy due to storage size limitations. To permit increased storage sizes, the Helm chart now includes new `resources` values under the `storage` component, as well as under `cven` and `kubelink` (note that amounts are in MiB not MB, and GiB not GB):

```
kubelink:
  resources:
    limits:
      memory: 500Mi
    requests:
      memory: 200Mi
      cpu: 200m

cven:
  resources:
    limits:
      memory: 300Mi
    requests:
      memory: 100Mi
      cpu: 250m

storage:
  resources:
    limits:
      memory: 500Mi
    requests:
      memory: 200Mi
      cpu: 100m
```

- **Pod to pod flows and pod labels are missing from Explorer search results** (E-116271, E-116272)
  In CLAS-enabled clusters, Explorer was not showing pod labels, only workload labels. In addition, Explorer did not return some traffic flows, even when trying with label-based search, or port-based search, or even searching using workload labels + pod labels. Also, pod traffic was being mapped to workloads.

Updates for Core for Kubernetes 5.1.2

Kubelink

Resolved Issues

- **Helm Chart: etcd storage size limit** (E-115417)
  Kubelink in CLAS mode uses etcd as a local cache for policy and runtime data. The Helm Chart now accepts a new variable called `storage.sizeGi` to set the size (in GiB not GB) of ephemeral storage. The default value is 1.
- **Kubelink - Unable to process policy with custom iptables rules** (E-115250)
  Kubelink in CLAS mode failed to process policy received from the PCE when custom iptables rules were present, producing the error message "json: cannot unmarshal object into Go struct field."
- **Kubelink to PCE connectivity issues - connection reset by peer** (E-115049)
  CLAS-enabled Kubelink was entering degraded mode too soon because of PCE connectivity problems. Now Kubelink also retries requests after network and OS errors, which avoids premature degraded mode entry.
- **C-VEN reporting potentially blocked traffic between worker nodes** (E-114691)
  CLAS processing of outbound rules to a ClusterIP Service replaced the "All Services" destination in the rule with actual ports from the Kubernetes Service. If a destination label included a Kubernetes Service, this caused a missing iptables rule between nodes.
- **Max policy message size between Kubelink and C-VEN is too small** (E-113714)
  The default gRPC message size was set to too small of a value, which caused C-VENs to reject policy messages that were larger than this value. The default gRPC message size is now larger, to avoid this problem.

Updates for Core for Kubernetes 5.1.0

What's New in the 5.1.0 Release

The following are new and changed items in the 5.1.0 release from the previous releases of C-VEN and Kubelink:

- **New CLAS architecture option**
  Kubelink now can be deployed with a Cluster Local Actor Store (CLAS) module, which manages flows from C-VENs to PCE, and policies from PCE to C-VENs. The CLAS-enabled Kubelink tracks individual pods, and when they are created or destroyed, instead of this being communicated directly to the PCE. To migrate from an existing (non-CLAS) environment to a CLAS-enabled one, set the `clusterMode` parameter to `migrateLegacyToClas` in your deployment YAML file (typically named `illumio-values.yaml`). See the `README.md` file accompanying the Helm Chart for full details on this and other Helm Chart parameters.
- **Workloads more closely match Kubernetes architecture**
  In CLAS-enabled environments, workloads are now conceptually tied to their containers, instead of being referred to in context of their pods, which more closely matches Kubernetes practice. To reflect this change, such workloads in CLAS environments are called *Kubernetes Workloads*, regardless of what containers have been spun up or destroyed to run the applications. In non-CLAS environments, the existing term *Container Workloads* is still used as in prior releases, corresponding to Pods. In mixed environments (with both non-CLAS and CLAS-enabled clusters), the PCE UI shows both Container Workloads and Kubernetes Workloads, as appropriate.
- **Degraded mode for CLAS-enabled Kubelink**
  If a CLAS-enabled Kubelink detects that its connection with the PCE becomes unavailable (for example, due to connectivity problems or an upgrade), Kubelink by default enters a *degraded mode*. In this degraded mode, new Pods of existing Kubernetes Workloads get the latest policy version cached in CLAS storage. When Kubelink detects a new Kubernetes Workload with exactly the same label sets and in the same namespace as an existing Kubernetes Workload, Kubelink delivers the existing, cached policy to Pods to this new Workload. If Kubelink cannot find a cached policy (that is, when labels of a new Workload do not match those of any existing Workload in the same namespace), Kubelink delivers a “fail open” or “fail closed” policy based on the Helm Chart parameter `degradedModePolicyFail`. The degraded mode can also be turned on or off by the Helm Chart parameter `disableDegradedMode`.
- **Illumio annotations in CLAS mode specified on the workload and not on Pod's template**
  Illumio annotations when in CLAS mode are now specified on the Kubernetes Workload and not on the pod's template.

- **Docker support dropped**
  The Docker CRI is no longer supported as of the 5.0.0 release of Illumio Core for Kubernetes.

C-VEN

Resolved Issue

- **Permanently delete Kubernetes Workloads after certain period when they are unpaired** (E-112362)
  Kubernetes Workloads (from a CLAS environment) are pruned from the PCE one day (by default) after they are unpaired. The length of time that elapses (in seconds) before this pruning occurs is configurable with the `vacuum_entities_wait_before_vacuum_seconds` parameter, which is set in the PCE `agent.yml` file. The default value for this parameter is 86400 (24 hours).

Known Issues

- **When C-VEN starts first, a 404 from PCE when getting CLAS token** (E-109259)
  When C-VEN is started first, it tries to contact the PCE in order to obtain CLAS token, but receives a 404 error. This is expected behavior for this scenario, which is only momentary. Kubelink eventually starts normally, and C-VEN obtains the CLAS tokens as expected.
- **Helm install fails with Helm version 3.12.2 but works with 3.10** (E-108128)
  When installing with Helm version 3.12.2, the installation fails with a YAML parse error.
  Workaround: Use Helm version 3.10, or version 3.12.3 or later.
- **Re-adding node does not re-pair it** (E-98120)
  After deleting a node and re-adding the same node, the node does not reappear, and previously established policy disappears from the node.
  Workaround: Uninstall and re-install Illumio Core for Kubernetes from scratch with the node present.

Kubelink

Resolved Issues

- **CLAS: NodePort - pod rules are not removed after disabling rule** (E-111689)
  After disabling a NodePort rule that opens it to outside VMs, iptable entries for pods with a virtual service's targetPort were not being removed as expected. Now the pod no longer

remains opened. Host iptables are removed, so traffic does not go through, and the pod ports are properly closed.

- **CLAS - The etcd pod crashes when node reboots** (E-106236)
  The etcd pod would crash if one of the nodes in the cluster was rebooted.

Known Issues

- **CLAS-mode Kubelink pod gets restarted once when deploying Illumio Core for Kubernetes** (E-109284)
  The Kubelink pod is restarted after deploying Illumio Core for Kubernetes in CLAS mode. There is no workaround. Kubelink runs properly after this single restart.
- **CLAS: Container Workload Profile label change is not applied to Kubernetes Workloads, only to Virtual Services** (E-109168)
  When removing labels in a Container Workload Profile, existing Kubernetes Workloads that are managed by that profile do not have their labels changed automatically to labels based on annotations. These existing Kubernetes Workloads must be updated with the `kubectl apply` command for the labels change to take effect. New Kubernetes Workloads created after the profile label change will have the new labels.
  This works as designed.

Security Information for Core for Kubernetes 5.1

For information about security issues, security advisories, and other security guidance pertaining to this release, see Illumio’s Knowledge Base in Illumio's Support portal.