Release Notes

Published

2023-10-06

Cloud-Native Contrail Networking 23.3

Table of Contents

Introduction

Juniper Cloud-Native Contrail® Networking (CN2) is a cloud-native SDN solution that provides advanced networking capabilities to containerized cloud networking environments. CN2 is optimized for Kubernetes-orchestrated environments and can be used to connect, isolate, and secure cloud workloads and services seamlessly across private, public, and hybrid clouds.

These release notes accompany Release 23.2 of CN2. They describe new features, limitations, platform compatibility requirements, known behavior, and resolved issues in CN2.

See the Cloud-Native Contrail Networking (CN2) page for a complete list of all CN2 documentation.

What's New

IN THIS SECTION

Learn about new features introduced in CN2 Release 23.3.

CN2 on OpenShift

- Advanced Cluster Management (ACM) — Starting with Release 23.3, you have the option of installing CN2 using Advanced Cluster Management. ACM consists of a hub cluster that provides centralized control of managed clusters. After you set up the hub cluster, you use ACM to install or import the CN2 clusters that you want to manage.

  See Install Using Advanced Cluster Management.

- **Single Node OpenShift** — Starting with Release 23.3, CN2 can run on a single node OpenShift deployment. A single node OpenShift deployment consists of a single node that runs both the control plane and the workloads.

  See Install Single Node OpenShift.

- **Seamless User Management** — Starting with Release 23.3, you can configure CN2 to use the Dex OpenShift connector for authentication of CN2 Web UI users. With this option, OCP users can access the CN2 Web UI seamlessly without further configuration.

  See Configure User Management.

Advanced Virtual Networking

- **Subinterface Support with Multus**—Starting with Release 23.3, CN2 supports multiple network subinterfaces using the Multus "meta" plugin. The "meta" references the Multus multi-vendor support. To configure subinterfaces on pods, use the network definition tags `net.juniper.contrail.interfacegroup` and `net.juniper.contrail.vlan` in the annotations `cni-args` section of the YAML.

  See Subinterface Support with Multus.

- **Immutable IP Address**—Starting in CN2 Release 23.3, an immutable IP address for the vhost0 interface is supported. No user configuration is required for this behavior change.

Configure Services

- **Support Color Communities in CN2**—Starting with CN2 Release 23.3, BGP color extended communities are supported. CN2 supports configuration of a color extended community using color:0:<tag> or color:<tag> versus the hexadecimal value. Color communities are attached to the routes using routing policies.

  See Configure BGP Color Extended Communities.

Configure eBPF

- **eBPF Kernel Data Plane (Tech Preview)**—Starting in CN2 Release 23.3, CN2 supports an extended Berkeley Packet Filter (eBPF) data plane for the Linux kernel. An eBPF-based data plane enables programs to be loaded into the kernel for high-performance applications.

  See eBPF Kernel Data Plane (Tech Preview).

CN2 Security

- **Routing Policies**—Starting in CN2 Release 23.3, you can apply dynamic routing policies to network traffic. Routing Policies modify a route's path and attributes dynamically. With release 23.3, the manipulation and filtering of routes is more granular.

  See Routing Policies.

- **Enable Namespace Isolation by Default**—Starting in CN2 Release 23.3, a default tag for isolated namespaces is supported. With CN2, you can enable a cluster to create isolated namespaces by default.

  See Enable Namespace Isolation by Default.

- **Global Security Policy**—Starting in CN2 Release 23.3, the `selectors` field for global Contrail security policies is supported. The `selectors` field is a combination of the `podSelector` and `namespaceSelector` fields. Global Contrail security policies define allow and deny rules for ingress and egress traffic between workloads (pods) across clusters.

  See Global Security Policy.

What's Changed

IN THIS SECTION

Learn about what changed in this release for CN2 23.3.

API Server Reload Support

- Until Contrail Networking Release 23.2, you had to restart the contrail-api-server manually whenever its certificate is renewed for the new certificate to take effect.

  > **NOTE**: The deployment of contrail-api-server depends on Kubernetes contrail-api-tls certificate. And the contrail-api-server controller tracks the Kubernetes secret where its certificate is stored. If you revoke this certificate, the corresponding secret is removed, and contrail-api-server controller goes into a failed state.

  The cert-manager (an internal component) issues contrail-api-server a certificate. By default, all certificates are valid for 10 years and they are renewed within 15 days.

  Starting with Contrail Networking Release 23.3, you are not required to restart the contrail-api-server manually. With API Server Reload support, the contrail-api-server restarts automatically whenever the contrail-api-server's certificate is renewed.

Immutable IP Address

- **vhost0 is not assigned the IP address of the physical interface upon creation**— The MAC address of the physical interface remains, however. As a result, the vRouter agent still establishes a connection with the CN2 controller, but existing routes are not affected because the IP address of the interface is unchanged. In other words, all of host's established routes point to the physical fabric interface instead of vhost0. CN2 release 23.3 doesn't utilize vhost0 for host communication.

Tested Integrations

Starting in CN2 Release 23.1, Supported Platforms is now documented in CN2 Tested Integrations. This document includes integrations fully tested and validated by Juniper, including tested NICs and other software components.

Container Tags

Container tags are needed to identify the image files to download from the Contrail Container Registry during a Contrail Networking installation or upgrade.

The procedures to access the Contrail Container Registry are provided directly by Juniper Networks. The location of the files in the Contrail Container Registry changed for the CN2 software starting in Release 22.4. To obtain access credentials to the registry or if you have any questions about file locations within the registry, send an email to: contrail-registry@juniper.net.

The following table provides the container tag name for the image files for CN2 Release 23.3

Table 1: Container Tag—Release 23.3

Orchestrator Platform	Container Tag
• Kubernetes: 1.27.5, 1.26.5, 1.25.5 • Red Hat OpenShift: 4.12.13, 4.10.31, 4.8.39 • Amazon EKS: v1.24.10-eks-48e63af • RKE2 v1.27.1+rke2r1	23.3.0.180

Open Issues

IN THIS SECTION

Learn about open issues in this release for CN2 23.3.

General Routing

- CN2-3429: When fabric source NAT is enabled in an isolated namespace, traffic flows between pods in isolated namespaces and between pods in isolated and non-isolated namespaces.
  Workaround: Do not configure fabric source NAT on an isolated namespace.

General Features

- CN2-3256: cSRX workloads with sub-interfaces are not compatible with CN2.
- CN2-6327: When interface mirroring is enabled with the **juniperheader** option, only egress packets are mirrored.

  Workaround: Disable the **juniperheader** option to mirror both egress and ingress packets.
- CN2-5916: When 4 interfaces are configured in a bond interface on an X710 NIC, an mbuf leaf with traffic drop occurs.

  Workaround: Limit two interfaces in a bond configuration for an X710 NIC.
- CN2-10346: When restarting a vRouter pod on kernel-mode nodes where vhost0 is installed onto bond interfaces, the bond IP address is assigned to a bond secondary interface instead of a bond primary interface.

  Run the following script for the workaround:

  ```
  Bond-patch.txt
  text · 982 B

  #!/bin/bash

  set -x

  slave_list=($(ip addr show | grep SLAVE | awk '{ print $2 }' | sed 's/://'))Revision History
  for slave in "${slave_list[@]}"; do
    IFS=$' '
  ```

```
  bond=$(ip addr show dev ${slave} | grep SLAVE | awk -F'master ' '{print $2}' | awk -F'
' '{print $1}')
  IFS=$'\n'
  route_list=($(ip route show | grep ${slave}))
  for route in "${route_list[@]}"; do
    echo "route: ${route}"
    new_route=$(echo ${route} | sed "s/${slave}/${bond}/g")
    route_cmd=$(echo "ip route replace ${new_route}" | sed -e 's|["'\'']||g')
    eval ${route_cmd}
  done
  ipv4=$(ip addr show dev ${slave} | grep 'inet ' | awk '{ print $2 }')
  ipv6=$(ip addr show dev ${slave} | grep 'inet6 ' | awk '{ print $2 }')
  echo "slave: '${slave}', bond: '${bond}', ipv4: '${ipv4}', ipv6: '${ipv6}'"
  if [[ -n "$ipv4" ]]; then
    ip addr del ${ipv4} dev ${slave}
    ip addr add ${ipv4} dev ${bond}
  fi
  if [[ -n "$ipv6" ]]; then
    ip addr del ${ipv6} dev ${slave}
    ip addr add ${ipv6} dev ${bond}
  fi
```

- CN2-13314: The gateway service instance (GSI) does not work with a 4-byte ASN.

  Workaround: Use a 2-byte ASN when connecting workloads through the GSI service.

- CN2-17407: In compute nodes running the Intel N6000 SmartNIC with CN2 23.3, it is necessary to add 12 Byes to the actual MTU expected from the interface.

Red Hat OpenShift

- CN2-7787: The KubeVirt deployment in Openshift 4.10 fails intermittently.

  See Red Hat OCPBUGS-2535 for a workaround.

- CN2-13011: Red Hat OCP backup and restore fails.

  See Red Hat https://access.redhat.com/solutions/6964756 for a workaround.

- CN2-16593: Monitor API used to fetch Prometheus metadata fails on OCP.

  Some observability and monitoring widgets in the CN2 UI do not work in an OCP deployment.

The following are some of the widgets that might not render data in the UI:

- **Dashboard > Observability**: CN2 Workloads, Kubernetes Overview, Top Nodes by Workload Utilization, CN2 Overview-Workload.
- **Monitoring > Orchestration > Ingress**: CPU Utilization, Memory utilization.
- **On Monitoring > Orchestration > DNS**: Core DNS Details, Queries handled by Cluster.
- **On Monitoring > CN2 > Controllers > Analytics**: CPU utilization of database engines, Memory utilization of database engines.
- **On Monitoring > CN2 > Metrics**: Only a few metrics are populating.

Workaround: Use the Analytics API to obtain the date for the affected widgets.

CN2 Apstra Integration

- CN2-13607: In a CN2 Apstra deployment, Apstra takes several minutes to create a virtual network under a scaled scenario.
- CN2-13428: VNI does not update in Apstra in an Intra-VN topology.

  In CN2 Apstra integrated environments, updating the VNI associated to a VN is not supported through Apstra.

  Workaround: If you need to update any VNI parameters, delete the VN and recreate it with new VNI parameters.

CN2 and Kubernetes

- CN2-4508: Contrail virtual network subnet created through NAD can not have user defined gateway.

  Workaround: None.

- CN2-4822: You can not configure BGPaaS objects on nodes that host the Contrail controller and worker nodes on same physical host.

  Workaround: None. Production deployments run the Kubernetes worker nodes and controller in different physical hosts.

- CN2-8728: When you deploy CN2 on AWS EC2 instances, running Kubernetes service traffic and Contrail datapath traffic on different interfaces is not supported.

  Workaround: Do not deploy Kubernetes and data traffic on the same interface in AWS.

- CN2-10351: KubeVirt v0.58.0 does not support imagePullSecret, required for pulling images from the secure registry: enterprise-hub.juniper.net/contrail-container-prod/.

  Following these steps for the workaround:

  1. Install Docker.

  2. Create a local insecure registry.

  3. Restart Docker.

  4. Download the required containers. The containers are located at Release Userspace CNI - dpdk vhostuser interface support Juniper/kubevirt. These containers are stored as Assets.

  5. Load the containers.

  6. Tag and push the containers to the new insecure registry.

  7. Download operator.yaml and cr.yaml.

  8. Modify the kubevirt-operator.yaml to use your insecure registry.

- CN2-14895: Pods are being deployed more than the VMI capacity of the nodes.

  When a custom pod scheduler is configured with maximum VMI capacity as thresholds, if the pods are scheduled back-to-back in quick succession, it is possible that more pods are deployed than the configured threshold. This is due to the delay in data sync between the node and analytics.

  Workaround: Additional pod scheduling on the busy nodes will stop within a few seconds once the VMI data is synced between the nodes and analytics.

- CN2-15530: Packet loss is observed in CN2 flow stickiness when scaling up from one to many pods (non-ECMP to ECMP).

  During scale up flow stickiness is applicable only within the ECMP group. Scale up from one to many pods does not maintain flow stickiness.

  Workaround: Start with a minimum of 2 workloads and scale up.

- CN2-15461: BFD session is not coming up when healthcheck is associated with 2 BGPaaS objects.

  Workaround: In environments where BFD is used with BGPaaS, if firewall policy is configured, ensure that the policy rules allow port 4784 (BFD packets).

Security

- CN2-4642: In CN2, the network policy uses the reserved tags `application` and `namespace.` These tags conflict with Contrail's reserved resources.

  Workaround: Do not use the application and namespace labels to identify the pod and namespace resources.

- CN2-10012: If the network policy has a deny-all rule, removing it by updating the policy does not work.

  Workaround: Delete the policy and re-add it again.

CN2 Pipelines

- CN2-15876: Tests are triggered when files in a different folder from the one specified in the YAML file directory are committed. The `cn2networkconfig` folder is specified in the `values.yaml` as the directory for commits and files are merged tests expected to be triggered. Argo CD only supports syncing from the path specified in the Helm chart as a part of CN2 pipeline startup.

  Workaround: Only commit to the `cn2networkconfig` directory.

- CN2-16034: Auto-created CN2 objects puts Argo out-of-sync after the commit. Creating a NAD starts the virtualRouter and subnets which are flagged as out-of-sync by Argo.

  Workaround: Add the `resource.exclusions`: in **charts/argo-cd/templates/argocd_sa.yaml**

  Workaround added to Helm chart:

```
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: argocd
  labels:
    app.kubernetes.io/name: argocd-cm
    app.kubernetes.io/part-of: argocd
  name: argocd-cm
data:
  resource.exclusions: |
    - apiGroups:
```

```
    - "*"
    kinds:
    - VirtualNetwork
    clusters:
    - "*"
timeout.reconciliation: 2s
```

Resolved Issues

You can research limitations that are resolved with this release at:

Resolved Issues in CN2 Release 23.3 .

Use your Juniper Support login credentials to view the list. If you do not have a Juniper Support account, you can register for one here.

Requesting Technical Support

IN THIS SECTION

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.

- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.

- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://supportportal.juniper.net/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://support.juniper.net/support/requesting-support/
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/

Revision History

- 29 September 2023—Revision 7
- 30 June 2023—Revision 6
- 30 March 2023—Revision 5
- 19 December 2022—Revision 4
- 23 September 2022—Revision 3
- 22 June 2022—Revision 2
- 02 May 2022—Revision 1, initial release