Extreme Visibility Manager Deployment Guide

6.0.0

9037083-00 Rev AA
June 2021

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: https://www.extremenetworks.com/support/policies/open-source-declaration/

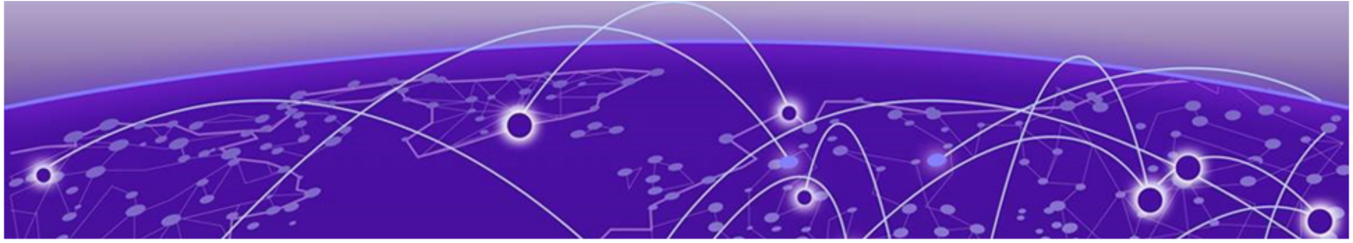

Table of Contents

Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as ExtremeSwitching switches or SLX routers, the product is referred to as *the switch* or *the router*.

Table 1: Notes and warnings

Icon	Notice type	Alerts you to...
💡	Tip	Helpful tips and notices for using the product.
📝	Note	Useful information or instructions.
➡	Important	Important features or instructions.

Table 1: Notes and warnings (continued)

Icon	Notice type	Alerts you to...
⚠	Caution	Risk of personal injury, system damage, or loss of data.
⚠	Warning	Risk of severe personal injury.

Table 2: Text

Convention	Description
`screen displays`	This typeface indicates command syntax, or represents information as it appears on the screen.
The words *enter* and *type*	When you see the word *enter* in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says *type*.
Key names	Key names are written in boldface, for example **Ctrl** or **Esc**. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press **Ctrl**+**Alt**+**Del**
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
italic text	Italic text indicates variable content.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation

Release Notes

Hardware and software compatibility for Extreme Networks products

Extreme Optics Compatibility

Other resources such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to The Hub.
2. In the list of categories, expand the **Product Announcements** list.

3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

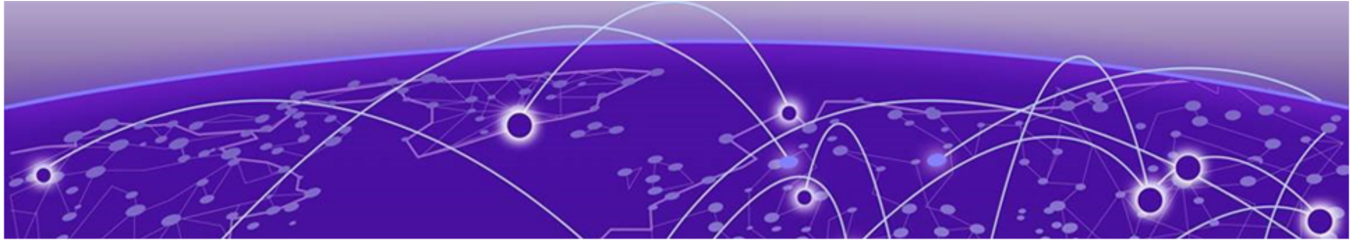

What's New In This Document

Extreme Visibility Manager 6.0.0 differs from previous releases in that it has a new architecture and a new user interface. This document is new for this release.

For information about the features in this release, see the *Extreme Visibility Manager Release Notes, 6.0.0*.

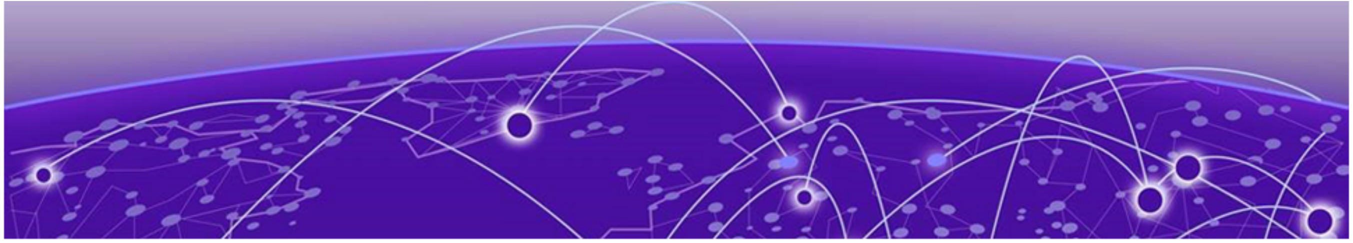

Preparing to Deploy Visibility Manager

This section describes the requirements and configuration that must be in place before you deploy Visibility Manager.

Extreme Visibility Manager Functions

Extreme Visibility Manager (Visibility Manager), a Kubernetes-based microservices application, provides centralized device and policy management as part of the Extreme Visibility solution.

Visibility Manager supports several network packet broker devices. Although devices have different functionality and different configuration methods, Visibility Manager seamlessly interacts with all supported devices for simplified management.

You use Visibility Manager to perform much of the same traffic configuration that you might otherwise perform from the command-line interface of your network packet broker operating system. And then you use Visibility Manager to analyze the traffic for insight into issues such as network usage, load-balancing irregularities, and security threats. For more information, see "Understanding Packet Broker Functions" in *Extreme Visibility Manager Administration and User Guide, 6.0.0*.

Visibility Manager managed objects work together to accomplish most packet broker functions. You configure these objects from the user interface.

Table 4: Managed objects

Object	Description
Ports and port channels	The interfaces on which traffic enters and exits the packet broker device. You can associate ports and port channels with ingress groups and egress.
Egress	A port or port channel that you associate with an egress policy, which identifies the actions to take on egress traffic.
Egress group	A set of interfaces and ports on which traffic is forwarded after a policy is applied.

Table 4: Managed objects (continued)

Object	Description
Ingress group	A collection of ports, port channels, and tunnels on which monitored traffic is received. You can select several actions to perform on the incoming traffic and you can associate the ingress group with an ingress policy.
Policy rule matches	The parts of a packet header that a rule targets, such as the source port or the payload length. One or more rules constitute a match. You associate matches with ingress or egress policies.
Ingress policy (or route map)	The actions to apply to inbound packets. You can associate policy rule matches and egress groups, and select other actions such as packet slicing and scope shift.
Egress policy (or listener policy)	The actions to apply to outbound packets. You can associate policy rule matches and select other actions such as packet slicing and header stripping.
User-defined access list (UDA)	The rules and matches created for or reconciled from MLX devices.
Transport tunnel termination and encapsulation	The GRE or ERSPAN tunnels to associate with ingress groups or egress.

Visibility Manager Deployment Model

With Visibility Manager, you can manage packet broker devices across multiple data centers and geographic locations.

You can separate your network into regions and further separate regions into zones. A zone is a set of locations (at least one and no more than five) where packet broker devices reside.

Note
See the *Extreme Visibility Manager Release Notes, 6.0.0* for any limitations on supported regions.

In this example, a map of the United States is separated into seven regions. Each region contains one or more zones. Each zone is a state with a data center.

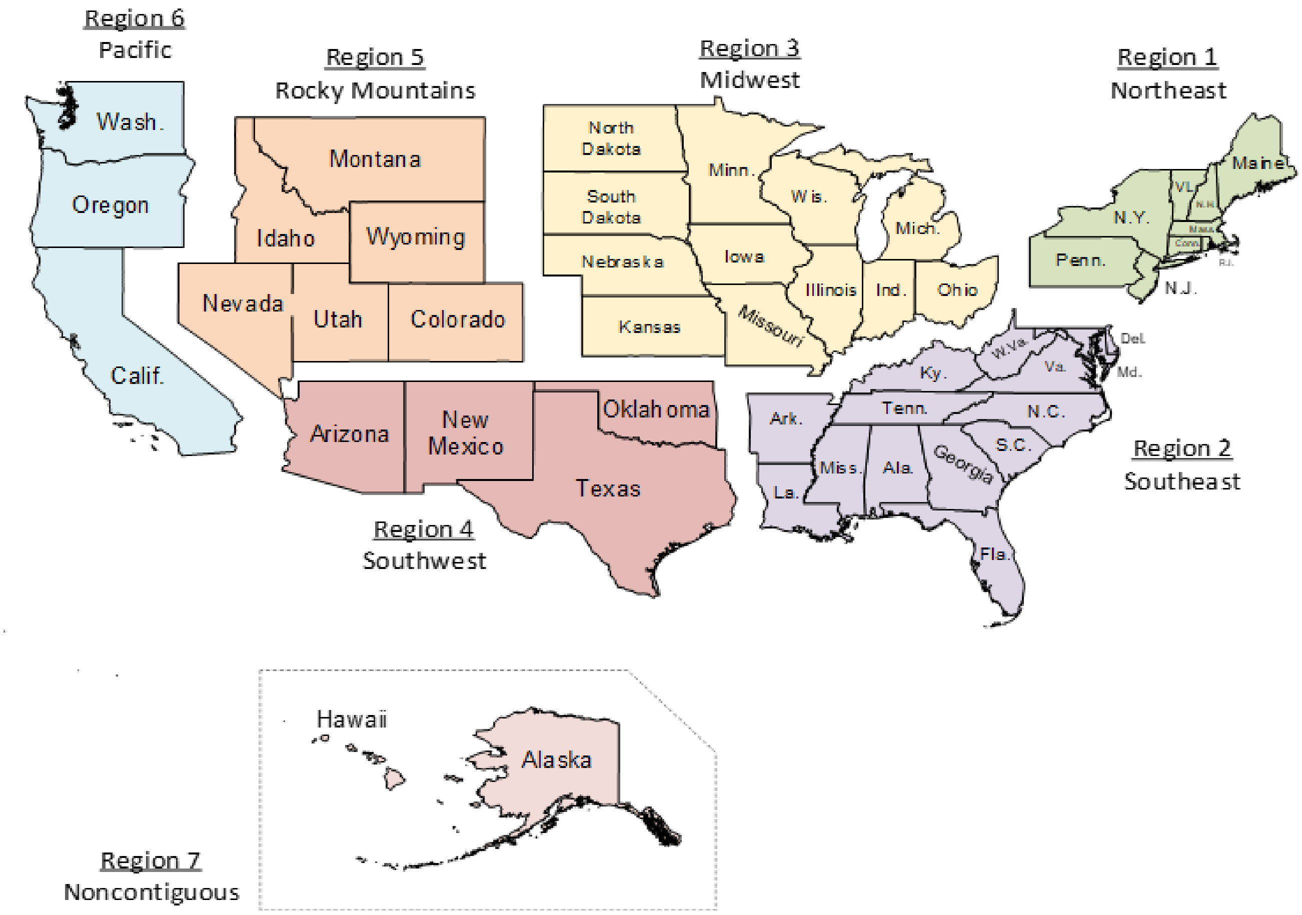

Figure 1: Regions and zones

Each region is a Kubernetes cluster that is managed by a control plane. Having multiple small clusters rather than one large cluster ensures better fault isolation. If one cluster fails, the other clusters provide failover assistance. All clusters are deployed on virtual machines (VMs).

In the following diagram, you can see the Kubernetes cluster, with the control plane, region nodes, and zone nodes.

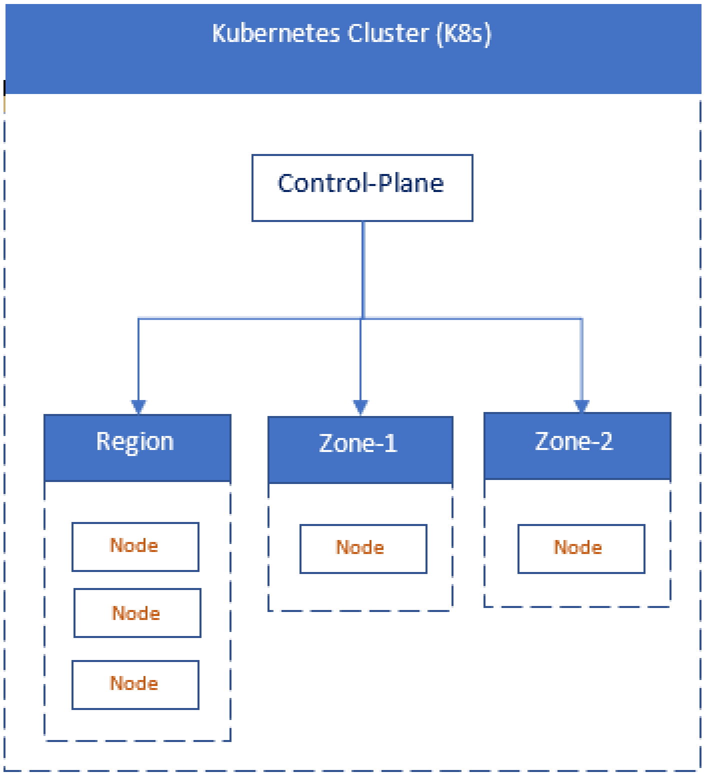

Figure 2: Cluster architecture

In this next diagram, you see how zones can have multiple locations, which tell you where packet broker devices are located.

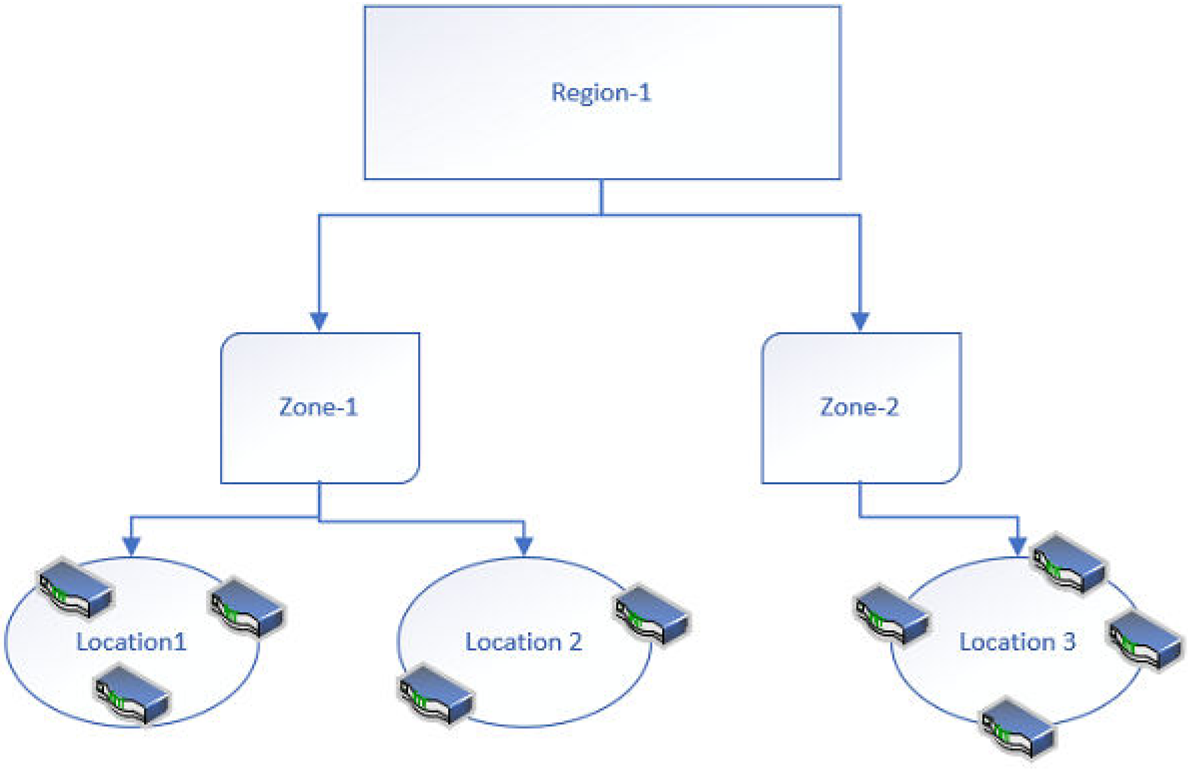

Figure 3: Zone architecture

You identify all zones and regions in a CSV file that is incorporated into Visibility Manager during installation. For more information, see Create a Location Definition File on page 15.

For information about deploying the control plan, region, and zone VMs, see Installing Visibility Manager on page 16.

Supported Devices

Extreme Visibility Manager supports several devices and their software.

Table 5: Supported devices and software

Device	Supported Software
Extreme 9920	Extreme 9920 software, version 21.1.0.0, with the NPB application
ExtremeRouting MLX series	NetIron 06.3.00d
ExtremeSwitching SLX 9140	SLX-OS 18s.1.03a, SLX-OS 18s.1.03b
ExtremeSwitching SLX 9240	SLX-OS 18s.1.03a, SLX-OS 18s.1.03b

System Requirements

Extreme Visibility Manager is installed on multiple virtual machines (VMs).

VM requirements

VM Type	Minimum Number of VMs	System Requirements	Maximum Devices
Control plane	1	• 2 vCPU • 4 GB RAM • 32 GB storage	N/A
Region	3	• 4 vCPU • 16 GB RAM • 200 GB storage	100 devices per region
Zone	1	• 2 vCPU • 4 GB RAM • 32 GB storage	25 devices per zone

System prerequisites

- Install `libguestfs-tools` (for quick installation) on the Hypervisor where VMs are hosted.

  ```
  sudo yum install libguestfs-tools
  ```
- Do not use the `192.168.0.0/16` series of IP addresses as management IP addresses for VMs.
- Do not use capital letters in host names.
- Ensure that all VMs have the same time zone as the devices that you want to monitor.

  ```
  timedatectl set-timezone <time-zone>
  ```

Supported connection protocols

Connections between Visibility Manager and the Extreme 9920 device are over secure TLS.

Connections between Visibility Manager and SLX or MLX devices are over UDP without TLS.

Browser requirements

You can access the Visibility Manager user interface with the following browsers:

- Google Chrome
- Mozilla Firefox

Certificate requirements

Visibility Manager uses HTTPS and requires self-signed certificates.

Create a Location Definition File

The Location Definition file (in CSV format) identifies regions and their associated zones and managed locations.

About This Task

You create the `locations.csv` file before you install Visibility Manager because you need to update the file during the installation process.

Procedure

1. Create a CSV file with a file name of `locations.csv`.
2. Add content to the file using the following format, allowing one row per location-name.

```
<country-name>,<region-name>,<zone-vm-ip>,<zone-name>,<zone-vm-ip>,<zone-host-name>,
<location-name>,<longitude>,<latitude>
```

For example:

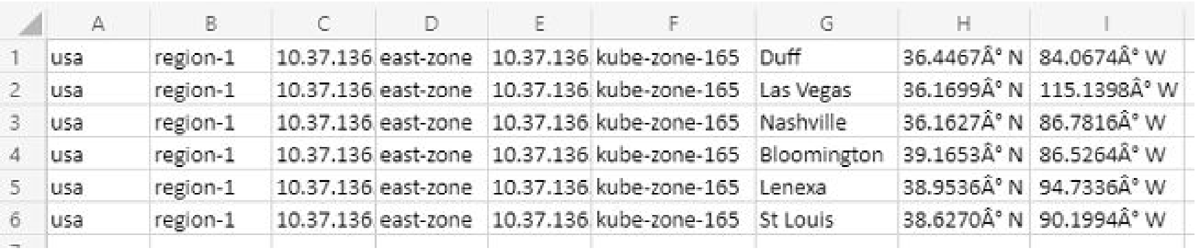

	A	B	C	D	E	F	G	H	I
1	usa	region-1	10.37.136	east-zone	10.37.136	kube-zone-165	Duff	36.4467Â° N	84.0674Â° W
2	usa	region-1	10.37.136	east-zone	10.37.136	kube-zone-165	Las Vegas	36.1699Â° N	115.1398Â° W
3	usa	region-1	10.37.136	east-zone	10.37.136	kube-zone-165	Nashville	36.1627Â° N	86.7816Â° W
4	usa	region-1	10.37.136	east-zone	10.37.136	kube-zone-165	Bloomington	39.1653Â° N	86.5264Â° W
5	usa	region-1	10.37.136	east-zone	10.37.136	kube-zone-165	Lenexa	38.9536Â° N	94.7336Â° W
6	usa	region-1	10.37.136	east-zone	10.37.136	kube-zone-165	St Louis	38.6270Â° N	90.1994Â° W

Figure 4: Sample locations.csv file

Note
Do not insert spaces after commas or in any fields other than the geographical location fields. Ensure that the IP addresses are valid. Zone names and host names can consist of numeric characters and lowercase alphabetic characters.

3. Save the file to a location that is accessible from the region VMs.

   You need to be able to edit the file from the region VMs during installation.

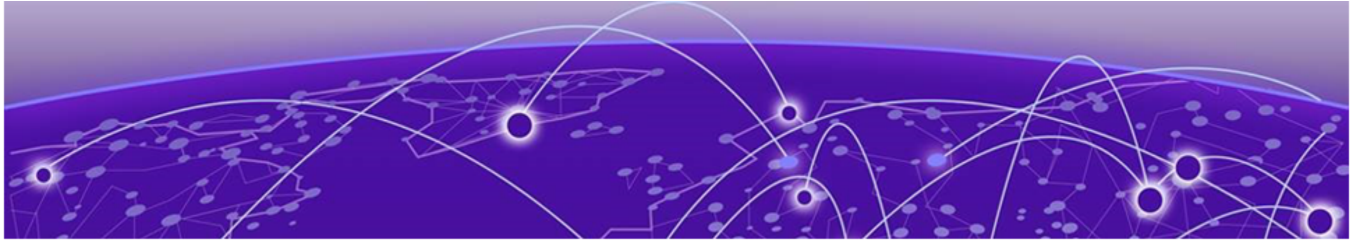

Installing Visibility Manager

This section describes the process for installing Visibility Manager VMs and then verifying the installation.

Visibility Manager VMs **must** be installed in the following order:

- Control plane VM
- Region 1 VM
- Region 2 and 3 VMs and the zone VM, which can be installed at the same time (in parallel)

Installation Worksheet

Gather the following information before you begin the installation process.

Table 6: Worksheet

Host	Host Name	Host IP Address	Subnet	DNS IP Address	Gateway
Control plane					
Region 1					
Region 2					
Region 3					
Zone					

Download and Untar the Build

This is part one of a five-part process.

Procedure

1. Download the `xvm-<version-build>.tar` file.

2. Untar the build directory.

   ```
   #tar -xvf xvm-<version-build>.tar
   ```

   Several files are untarred, including *.conf files and qcow2 files for each control plane, region, and zone VM.

   > Note
   > For multiple regions, make copies of the region qcow2 file.

3. Create different directories for each VM in which to store the qcow2 files and the `xvmconf` files.

   ```
   # cp xvm_<filename>/controlplane-xvm_<filename>.qcow2 controlplane1-xvm.qcow2

   # cp xvm_<filename>/region-xvm_<filename>.qcow2 region1-xvm.qcow2

   # cp xvm_<filename>/zone-xvm_<filename>.qcow2 zone1-xvm.qcow2
   ```

What to Do Next

Create the Control Plane Virtual Machine on page 17

Create the Control Plane Virtual Machine

This is part two of a five-part process.

Before You Begin

Download and Untar the Build on page 16

Procedure

1. Open the `controlplane.conf` file, which you obtained when you untarred the release tar file.
2. Update `controlplane.conf` as described in the comments of the file.

   > Note
   > In `controlplane.conf`, do not enclose values in double-quotes. For example, to configure NTP, NTP_CONF=1 is correct. NTP_CONF="1" is incorrect.

3. Copy `controlplane.conf` to `xvmconf`.

   ```
   # cp controlplane.conf xvmconf
   ```

4. Copy the `xvmconf` file to the qcow image.

   ```
   # virt-copy-in -a <absolute-path-for-controlplane-qcow2-file> xvmconf /etc/
   ```

5. Install the control plane VM.

   For example:

   ```
   virt-install --name controlplane198  --ram 4096 --vcpus 2
   --disk path=/home/user/sbrBuilds/controlplane-xvm.qcow2,bus=sata  --nographics --
   import
   --os-variant centos7.0 --network bridge:<name-of-bridge> --console
   pty,target_type=virtio
   ```

   Note

   The *<name-of-bridge>* variable is the name of the bridge that is built into your host system. You can use the **ifconfig** command to identify the bridge configuration of the host.

   The VM starts up.

6. Log in with the following credentials.
   - user name: root
   - password: password
7. Select **-network** (according to your Hypervisor setup) and follow the automated instructions to complete the installation.

What to Do Next

Create the Region Virtual Machines on page 18

Create the Region Virtual Machines

This is part three of a five-part process.

Before You Begin

Create the Control Plane Virtual Machine on page 17

About This Task

Tip

After you create the first of the three region VMs, you can create the remaining region VMs and the zone VM at the same time.

Procedure

1. Open the `region.conf` file, which you obtained when you untarred the release tar file.
2. Update `region.conf` for one region as described in the comments of the file.

   Note

   In `region.conf`, do not enclose values in double-quotes. For example, to configure NTP, NTP_CONF=1 is correct. NTP_CONF="1" is incorrect.

3. Copy `region.conf` to `xvmconf`.

   ```
   # cp region.conf xvmconf
   ```

4. Copy the `xvmconf` file to the qcow image.

   ```
   # virt-copy-in -a <absolute-path-for-region-qcow2-file> xvmconf /etc/
   ```

5. Repeat steps 2 through 4 once for each region you are creating.
6. Install the region VMs.

```
virt-install --name region199  --ram 16384 --vcpus 4
--disk path=/home/user/sbrBuilds/region1-xvm.qcow2,bus=sata  --nographics --import
--os-variant centos7.0 --network bridge:<name-of-bridge> --console
pty,target_type=virtio

virt-install --name region200  --ram 16384 --vcpus 4
--disk path=/home/user/sbrBuilds/region2-xvm.qcow2,bus=sata  --nographics --import
--os-variant centos7.0 --network bridge:<name-of-bridge> --console
pty,target_type=virtio

virt-install --name region201  --ram 16384 --vcpus 4
--disk path=/home/user/sbrBuilds/region3-xvm.qcow2,bus=sata  --nographics --import
--os-variant centos7.0 --network bridge:<name-of-bridge> --console
pty,target_type=virtio
```

> Note
> The `<name-of-bridge>` variable is the name of the bridge that is built into your host system. You can use the **`ifconfig`** command to identify the bridge configuration of the host.

The VMs start up.

7. Log in with the following credentials.
   - user name: root
   - password: password
8. Follow the automated instructions to complete the installation.

What to Do Next

Create the Zone Virtual Machine on page 19

Create the Zone Virtual Machine

This is part four of a five-part process.

Before You Begin

Create the Region Virtual Machines on page 18

About This Task

All alpha characters in the host name of the zone node must be lowercase.

Procedure

1. Open the `zone.conf` file, which you obtained when you untarred the release tar file.
2. Update `zone.conf` as described in the comments in the file.

> Note
> In `zone.conf`, do not enclose values in double-quotes. For example, to configure NTP, NTP_CONF=1 is correct. NTP_CONF="1" is incorrect.

3. Copy `zone.conf` to `xvmconf`.

```
# cp zone.conf xvmconf
```

4. Copy the `xvmconf` file to the qcow image.

   ```
   # virt-copy-in -a <absolute-path-for-zone-qcow2-file> xvmconf /etc/
   ```

5. Install the zone VM.

   ```
   virt-install --name zone201  --ram 4096 --vcpus 2
   --disk path=/home/user/sbrBuilds/zone1-xvm.qcow2,bus=sata  --nographics --import
   --os-variant centos7.0 --network bridge:<name-of-bridge> --console
   pty,target_type=virtio
   ```

   > Note
   >
   > The *<name-of-bridge>* variable is the name of the bridge that is built into your host system. You can use the **ifconfig** command to identify the bridge configuration of the host.

   The VM starts up.

6. Provide the following credentials to log in.
   - user name: root
   - password: password
7. Follow the automated instructions to complete the installation.

What to Do Next

Configure and Verify the System on page 20

Configure and Verify the System

This is part five of a five-part process.

Before You Begin

Create the control plane, region, and zone VMs, and the locations.csv file.

- Create the Control Plane Virtual Machine on page 17
- Create the Region Virtual Machines on page 18
- Create the Zone Virtual Machine on page 19
- Create a Location Definition File on page 15

Procedure

1. Ensure that the installations on the control plane, region, and zone VMs (nodes) are complete.
2. On the control plane VM, verify that `etcd` and `haproxy` are running.

   ```
   # systemctl status haproxy
   # systemctl status etcd3
   ```

3. On each region VM, verify that the `patroni` service is running.

   ```
   # systemctl status patroni
   ```

4. On any region VM, determine whether leader election is complete.

   ```
   # patronictl -c /opt/app/patroni/etc/postgresql.yml list postgres
   ```

   Output of the command identifies the IP address of the host with the Leader role.

5. On the control plane VM, create the join token.

```
# kubeadm token create --print-join-command

kubeadm join 10.37.138.217:6443 --token cmtjhj.fah33qgst7gl0z7w
--discovery-token-ca-cert-hash
sha256:9de0fcb3e7c5a5aa6a3ecbea0b9e9c3b3c187b6777b5c2746b2ce031240875be
#
```

6. On each region and zone VM, run the highlighted portion of the output in step 5.
7. On the control plane VM, verify that the nodes are joined.

```
# kubectl get nodes

NAME       STATUS     ROLES                  AGE     VERSION
control    Ready      control-plane          6m50s   v1.20.5
regions1   Ready      <none>                 58s     v1.20.5
zones1     Ready      <none>                 53s     v1.20.5
#
```

The wait time for VMs to join is typically 1 to 2 minutes. In the output of the command, regions and zones should be in Ready state.

8. On each region and zone VM, verify that Docker images are loaded.

```
# docker images
```

9. On all VMs, add default DNS.
   a. On the control plane VM, determine the cluster IP address.

```
# kubectl get svc -n kube-system

NAME    TYPE    CLUSTER-IP   EXTERNAl-IP   PORTS            AGE
kube-dns   ClusterIP 10.x.x.x     <none>      53/UDP,53/TCP,9153/TCP   13m
```

   b. Add the cluster IP address to `/etc/resolv.conf`.

```
# echo "nameserver 10.x.x.x" > /etc/resolv.conf
# cat /etc/resolv.conf
nameserver 10.x.x.x
```

10. Verify that all system pods are in Running state.

```
# kubectl get pods -A
```

The Status column in the output shows the state.

11. If the output of step 10 shows that coredns pods are in CrashLoopBackOff state, take the following steps. Otherwise, skip to step 12 on page 22.
    a. Run the **`kubectl edit cm coredns -n kube-system`** command.

```
}
     ready
     kubernetes clusster.local in-addr.arpa ip6.arpa {
        pods insecure
        fallthrough in-addr.arpa ip6.arpa
        ttl 30
     }
     prometheus :9153
     forward . /etc/resolv.conf {
        max_concurrent 1000
     }
     cache 30
     loop
     reload
     loadbalance
  }
kind: ConfigMap
```

```
metadata:
  creationTimestamp: "2021-03-20T06:04:50Z"
  name: coredns
  namespace: kube-system
  resourceVersion: "274"
  uid: 89a2b293-c0a3-b4b8-f262834020b
```

b. Delete the loop, save, and then exit.

c. Delete both coredns pods.

```
kubectl delete pod coredns-<full-pod-name> -n kube-system
```

d. Repeat step 10 on page 21 to ensure that the coredns pods are in Running state.

12. Set node labels for each region and zone using the **`kubectl label nodes <hostname><label>`** syntax, where <label> is in the following format: reg1 for regions 1, 2, and 3, and reg1-zone1, reg1-zone2 for different zones in a region.

```
kubectl label nodes cat-region1-evm region=reg1
kubectl label nodes cat-region2-evm region=reg1
kubectl label nodes cat-region3-evm region=reg1
kubectl label nodes cat-zone1-evm zone=reg1-zone1
```

The final command differs from the others in two ways. It has a different label and the label value is assigned.

13. Copy the `locations.csv` file from the qcow build directory to one of the region VMs.

14. Edit the `locations.csv` file to reflect zone VM configuration.

Do not insert spaces after commas or in any fields other than the geographical location fields. Ensure that the IP addresses are valid. Zone names and host names can consist of numeric characters and lowercase alphabetic characters.

```
usa,region-1,10.37.138.187,east-zone,10.37.138.187,zone-187,Duff,
36.4467° N,84.0674° W
usa,region-1,10.37.138.188,west-zone,10.37.138.188,zone-188,Las Vegas,
36.1699° N,115.1398° W
```

This example shows two zones in one region, with IP addresses of 10.37.138.187 and 10.37.138.188 and host names of zone-187 and zone-188. You must edit the highlighted fields.

15. On each region VM, create the following directory and copy the `locations.csv` file into the directory.

```
mkdir -p /opt/crms
cp locations.csv /opt/crms
```

16. On the control plane VM, perform the following.

```
cd /etc/xvm/controlplane_node_binaries
./loadPodsInControlPlaneNode.sh
```

The wait time for the script to finish running is typically 8 to 10 minutes.

17. Verify that pods are in Running state.

```
# kubectl get pods -n xvm
```

Installation Troubleshooting

This topic provides solutions for some installation issues.

Services in Pending state

Issue: Services are in Pending state after you run **`kubectl get pods -n xvm`**.

Reason: Labels for the region and zone nodes are incorrectly configured.

Solution:

1. Check the label configuration: **`kubectl get nodes --show-labels`**.
2. Remove the incorrect label: **`kubectl label node <node-name> <label-name>-`**.
3. Configure the correct label: **`kubectl label node <node-name> <label-name>=<value>`**.

   For example: **`kubectl label node zonal1-node-xvm region=reg1`**.
4. On the control plane VM, perform the following:
   ```
   cd /etc/xvm/controlplane_node_binaries
   ./loadPodsInControlPlaneNode.sh
   ```

Services in ErrImagePull or ImagePullBackOff state

Issue: Services are in ErrImagePull or ImagePullBackOff state after you run **`kubectl get pods -n xvm`**.

Reason: You tried to create a pod that references an image name or tag that does not exist. This issue can be caused by a version mismatch or by an incorrect label configuration.

Solution for a version mismatch:

1. Determine whether the Docker image exists on the region and zone nodes: **`docker images`**.
2. Load the same version of Visibility Manager on all nodes in the cluster.
3. Load the correct Docker image.
   ```
   docker load < /etc/xvm/<ABC>_node_binaries/<docker_image.tar.gz>
   ```

Solution for an incorrect label configuration:

1. Check the label configuration: **`kubectl get nodes --show-labels`**.
2. Remove the incorrect label: **`kubectl label node <node-name> <label-name>-`**.
3. Configure the correct label: **`kubectl label node <node-name> <label-name>=<value>`**.

   For example: **`kubectl label node zonal1-node-xvm region=reg1`**.
4. On the control plane VM, perform the following:
   ```
   cd /etc/xvm/controlplane_node_binaries
   ./loadPodsInControlPlaneNode.sh
   ```

crms-ms does not start and pod is in CrashLoopBackOff state

Issue: The crms-ms does not start or does not run correctly and pods are in CrashLoopBackOff state.

Reason: An incorrect parameter may be configured in the `/opt/crms/locations.csv` file.

Solution:

1. Correct the locations.csv file on all region nodes as described in Configure and Verify the System on page 20.

Do not insert spaces after commas or in any fields other than the geographical location fields. Ensure that the IP addresses are valid. Zone names and host names can consist of numeric characters and lowercase alphabetic characters. For example:

```
usa,region-1,10.37.138.187,east-zone,10.37.138.187,zone-187,Duff,
36.4467° N,84.0674° W
usa,region-1,10.37.138.188,west-zone,10.37.138.187,zone-188,Las Vegas,
36.1699° N,115.1398° W
```

2. On the control plane VM, perform the following:

```
cd /etc/xvm/controlplane_node_binaries
./loadPodsInControlPlaneNode.sh
```

Nodes not listed after you run **kubectl get nodes**

Issue: Nodes are not listed after you run **`kubectl get nodes`** on the control plane VM.

Reason: There are two possibilities. An incorrect join token was applied to the region and zone VMs. Or the host name is incorrectly configured on the region and zone VMs. For example, the same host name is used for more than one node in the `xvmconf` file.

Solution for an incorrect join token:

1. On the control plane VM, create the join token:

```
kubeadm token create --print-join-command
```

.

2. On the nodes that were not listed, remove any existing tokens: **`kubeadm reset`**.
3. Copy the join token on the nodes that were not listed: **`kubeadm join`**.

   For more information, see Configure and Verify the System on page 20.
4. Continue with the Configure and Verify the System on page 20 process after copying the join token.

Solution for an incorrect host name:

1. Verify the host name on all nodes in the cluster: `cat /etc/hosts/`.
2. Correct the xvmconf file in this location: `vi /etc/xvmconf`.
3. Rerun the startup script:

```
/etc/xvm/<node>_binaries/<node>_startup.sh
```

4. Set node labels for each region and zone using the **`kubectl label nodes <hostname><label>`** syntax, where <label> is in the following format: reg1 for regions 1, 2, and 3, and reg1-zone1, reg1-zone2 for different zones in a region. For example:

```
kubectl label nodes cat-region1-evm region=reg1
kubectl label nodes cat-region2-evm region=reg1
kubectl label nodes cat-region3-evm region=reg1
kubectl label nodes cat-zone1-evm zone=reg1-zone1
```

   The final command differs from the others in two ways. It has a different label and the label value is assigned.
5. On the control plane VM, perform the following:

```
cd /etc/xvm/controlplane_node_binaries
./loadPodsInControlPlaneNode.sh
```

Cluster node reconfiguration

You can reconfigure any of the region or zone nodes for reasons such as the following:

- Change the management IP address
- Change the host name
- Change any parameters in the xvmconf file

Take the following steps:

1. As needed, update the configuration parameters of the xvmconf file at `vi /etc/xvmconf`.
2. To delete a region or zone node from the control plane, take the following steps:,

    a. Run the following on the control plane node:

    ```
    kubectl cordon <node-name>
    kubectl drain <node-name> --ignore-daemonsets
    kubectl delete node <node-name>
    ```

    b. On the region or zone node, run **kubeadm reset**.

    The **drain** waits for graceful termination. Do not operate the node's VM until the command completes. To put the node back into service, run **kubectl uncordon <node-name>**, which makes the node able to be scheduled.
3. From the node where you updated xvmconf, run the startup script:

    ```
    /etc/xvm/<ABC>_node_binaries/<ABC>_startup.sh
    ```

    <ABC> can be a region or a zone. For example: **/region_startup.sh**.
4. On the control plane VM, create a new join token:

    ```
    kubeadm token create --print-join-command
    ```

    Alternatively, you can reuse a previously used join token if it is available.
5. Copy the join token on the region or zone node: **kubeadm join**.
6. Verify that the nodes have joined the cluster: **kubectl get nodes**.
7. Set node labels for each region and zone using the **kubectl label nodes <hostname><label>** syntax, where <label> is in the following format: reg1 for regions 1, 2, and 3, and reg1-zone1, reg1-zone2 for different zones in a region. For example:

    ```
    kubectl label nodes cat-region1-evm region=reg1
    kubectl label nodes cat-region2-evm region=reg1
    kubectl label nodes cat-region3-evm region=reg1
    kubectl label nodes cat-zone1-evm zone=reg1-zone1
    ```

    The final command differs from the others in two ways. It has a different label and the label value is assigned.
8. On any region node, verify that leader election is complete:

    ```
    # patronictl -c /opt/app/patroni/etc/postgresql.yml list postgres
    ```

    The output identifies the node with the Leader role and displays an * in the Pending Restart column.
9. On the leader region node, restart the `patroni` service:

    ```
    sudo systemctl restart patroni
    ```

10. On any region node, ensure nothing is in Pending Restart state:

    ```
    patronictl -c /opt/app/patroni/etc/postgresql.yml list postgres
    ```

11. On the control plane VM, perform the following:

    ```
    cd /etc/xvm/controlplane_node_binaries
    ./loadPodsInControlPlaneNode.sh
    ```