b. Enter the email associated with a vCommander user account in the text field. When a directory service has been configured, click the ellipsis to search for the user account.

c. Click **OK** when complete.

d. Repeat the sequence as required.

Troubleshooting

On certain occasions, information in a provisioning email may differ from the actual state of a service request. This can result from actions occurring after the email is sent or before it is received.

- A completion email may show a cost of zero even though the cost has been calculated for the service and is displayed on the Service Request Post-deployment landing page. In this case, use the hyperlink in the email to confirm the details.
- The first email sent about post-deployment may not contain the correct resource information that was part of the request or was modified as part of the deployment. In this case, add a "Wait for Service to Obtain IP Address and DNS name" step in a completion workflow to allow time for the resources to be determined properly. See Workflow Steps Reference.

Networking and IP Management

This article provides a general overview of Embotics® vCommander® networking.

See also the following articles in the Networking and IP Management category:

Configuring Networks

Configuring and Managing IP Pools

Network Fencing

Integrating BlueCat™ IP Address Management with vCommander

See also the following Knowledge Base articles:

Why Am I Getting an Automatic Private IP Addressing (169.x.x.x) Address?

Why are VMs Deployed with Manual MACs When the Base Image is Automatic?

Network configuration and automated deployment

When manually provisioning a VM (by manually deploying a service request, deploying a template, or cloning), the vCommander administrator has full control over network configuration for the deployed VM.

When you've set up automated deployment, how you configure networking for deployed VMs depends on how knowledgeable your users are and how much control you want to allow them. For example, if you're running a Dev/Test shop, you'd likely want to allow your users to choose the network zone and add adapters as required when requesting a service. You might also want to allow Dev/Test users to reconfigure networking settings for existing VMs. To satisfy these requirements, you would:

- assign network zones to networks

- add the Network element to the request form, allowing users to choose the network zone and add adapters as required
- set up an approval workflow to enable automated deployment for approved requests
- create a deployment destination for each of your user groups, with basic networking configuration

Automated deployment then assigns a network configured in the deployment destination that matches the network zone chosen on the request form.

If you're a service provider, on the other hand, your users likely don't know whether the VM they're requesting should be deployed into the DMZ or the Production zone, or whether they need extra network adapters. Instead, you may want to allow requesters to indicate their service requirements by selecting from a drop-down list (such as None, Backup or Monitoring), and then set up post-deployment configuration to handle these requirements. To set up a solution for this scenario, you would:

- add a list-type custom attribute to the request form, allowing users to specify what type of services they require
- set up an approval workflow to enable automated deployment for approved requests
- create a deployment destination for each of your user groups, with basic networking configuration
- set up a component-level completion workflow with a conditional step to configure networking based on information specified on the request form

Automated deployment then configures the new VM with the same NICs as the source template; the completion workflow adds a NIC and connects it to the appropriate network.

Viewing network details

For on-premise managed systems, vCommander provides a Networks tab at the following levels of the Operational tree and the VMs and Templates tree:

- Managed System
- Datacenter
- Cluster
- Host

SCVMM: Only the logical networks are displayed on the Networks tab. VM and virtual networks are not shown, and you can't assign zones to them.

See Network Properties for a description of the properties shown on the Networks tab.

For public cloud managed systems, vCommander provides a Subnets tab at the following levels of the Operational tree and the VMs and Templates tree:

- Managed System
- Region
- Availability zone or affinity group
- Virtual Private Cloud or Virtual Network
- Database (for AWS only)

See Subnet Properties for a description of the properties shown on the Subnets tab.

VM Network Assignment

You can view the network assignment for individual VMs by adding the Network property to a VM table or to the Details pane in the VM Summary tab. The Network property displays:

- For vCenter, the network name
- For AWS, the VPC name if the VM is in a VPC; otherwise, the value is blank
- For Microsoft Azure, the subnet name

This property is not displayed for SCVMM VMs.

You can also search for VM network assignment by going to **Tools > Search** and filtering by **Network** (in the Resources - Network category).

Order of precedence for network selection

The order of precedence for network selection for new VMs is as follows:

- Network deployment parameters specified by a request approver, or in a workflow step
- Network matching the network zone specified on the request form
- Network configured in the deployment destination wizard. When multiple networks are configured in the deployment destination, the first (alphabetically) is used.
- Network of source NIC

Configuring network resources in vCommander

vCommander provides several ways to configure networking.

Allowing users to specify networking information when requesting a service

Depending on how you configure the request form, when you add the Network form element, users requesting a VM can:

- change the network automatically selected by vCommander based on the deployment destination settings, by selecting a network zone
- add network adapters to a requested VM component

The network zones selectable on the form must match those configured for networks added to the deployment destinations available to the requesting user. If the user selects a network zone that isn't available on the target destination, automated deployment will fail.

To learn how to add the Network element to the new service request form, see:

- Adding a vCenter Service to the Catalog
- Adding an SCVMM Service to the Catalog
- Adding an AWS Service to the Catalog
- Adding an Azure Service to the Catalog

Note that you need to tag your networks with zones.

Assigning networks to automated deployment destinations

When you configure destinations for automated deployment placement, you configure network assignment for new VMs. You can add multiple networks to a destination.

If multiple networks of the same network zone are valid for a user, the first network is selected (alphabetically, case insensitive) for automated deployment.

Assigning networks after provisioning using completion workflows

You can also configure network assignment in a VM completion workflow. You can configure network assignment in the guest OS and at the hypervisor level.

Assigning networks during the request approval process

Use the `$NETWORK<x>=<network>$` deployment parameter to configure the network during the request approval process. Note that this parameter assigns the actual network, not the network zone. Request approvers can enter this parameter when approving a request; this parameter can also be used in scripts.

Assigning networks during manual deployment

Administrators can configure networking during manual provisioning.

Reconfiguring network resources

vCommander users can reconfigure network resources for deployed VMs.

Note for CentOS

Some versions of CentOS are not compatible with network customization when the base image includes a network interface. You can resolve this issue by removing the network interface from the base image. For more information, contact support@embotics.com.

Assigning Network Zones

A network zone is just a simple label that describe a network's purpose within your infrastructure — for example, DMZ, Corporate Intranet, Storage or Production. Network zones:

- enable your users to select a network when requesting a new service, without requiring an understanding of the underlying infrastructure. Requesters can also add or change network adapters for the service, and each can reside on a different network.
- help administrators keep track of a large list of networks
- simplify automated deployment in multi-tenant environments. For example, you can assign the same zone to different networks, create one service catalog entry for all organizations, and allow a user to select from several network zones. You can create a separate deployment destination for each organization and configure different networks in each destination.
- simplify automated deployment when using multiple managed system types. You can assign the same zone to networks on different managed systems, so that for example, your users just need to select "DMZ" whether they're requesting an AWS service or a vCenter service.

Important: Make sure networks are always available to back the network zones. If no network of the specified zone is available on the target managed system, automated deployment will fail.

See also Networking and IP Management for a general overview of Embotics® vCommander® networking.

Creating and managing network zones

Access through:	**Configuration menu > IP Pools and Networking > Networking tab**
Available to:	**vCommander Role of Superuser or Enterprise Admin**

You manage network zones on the Network Zones page. Three zones are created by default: Dev, DMZ and Production.

To create a zone: Click **Add**, enter a name and click **OK**.

To rename a zone: Select it in the list and click **Rename**.

To delete a zone: Select it in the list and click **Delete**. Note that you can't delete a zone that is currently applied to a network.

Applying a zone to a network

Access through:	**Views Menu > Operational or VMs and Templates View**
Available to:	**Administrator Access Rights**

To apply a zone to a network:

1. Navigate to a level of the tree in the Operational or VMs and Templates where the Networks tab or the Subnets tab appears. See Viewing Network Details above to learn more.
2. Select the **Networks** or **Subnets** tab.
3. Select one or more networks or subnets in the list.
4. Right-click and select **Set Network Zone**.
5. In the Set Network Zone dialog, select a network zone in the drop-down list and click **OK**.

   To remove a network zone assignment, select the blank entry in the drop-down list.

Configuring and Managing IP Pools

IP pools are used by Embotics® vCommander® administrators to define ranges of IP addresses to be assigned as vCenter VMs are deployed on a specified network. Typically, IP pools would be used to reserve resources for particular users or groups or to make sure that certain ranges are used for particular purposes.

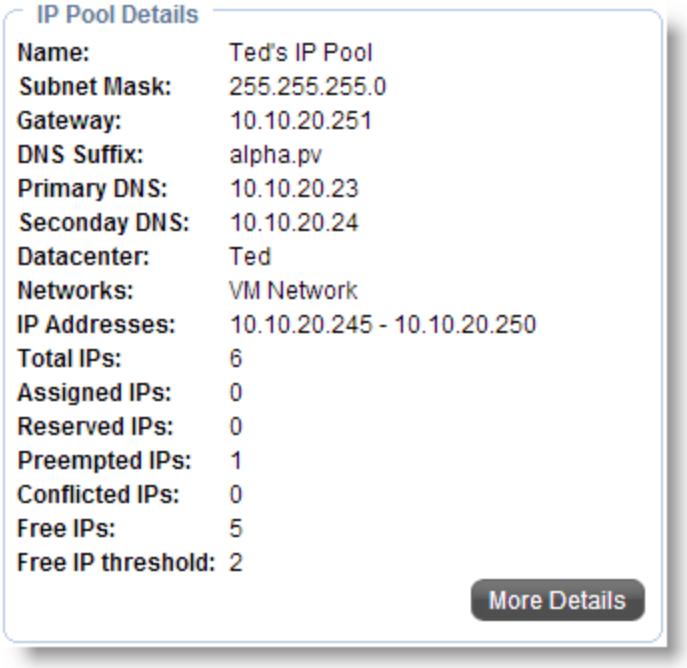

IP Pool Details

It's important to understand that IP pools do nothing to actually reserve the IP addresses or prevent them from being statically assigned to VMs by users. However, you can configure the system to send you an email when an IP conflict has been detected.

There are several steps to configuring IP pools:

1. Add an IP Pool.
2. Create customization specifications.
3. Assign the customization specifications to your service catalog entries.
4. Update the automated deployment placement.
5. Configure email notification.

vCommander IP pools are supported for vCenter only. See Integrating BlueCat™ IP Address Management with vCommander to learn how to create host records or DHCP reservations from BlueCat™ IPAM during provisioning with vCommander.

See also Preparing vCenter Networking for Fencing to learn how to configure network fencing — isolated VM configurations that avoid IP or MAC address conflicts, but have full network access. Network fencing allows you to work with multiple live instances of the same configuration on the same network.

Add an IP pool

Access through:	**Configuration menu > IP Pools and Networking > IP Pools tab**
Available to:	**vCommander Roles of Superuser and Enterprise Administrator**

Before adding an IP pool, take whatever steps are necessary to make sure users will not try to assign the IPs manually to VMs or real computers on the network. For example, if the range you want to use for your pool is included in your DHCP scope, update or split the scope so the addresses are no longer included.

1. On the IP Pools tab, click **Add**.
2. On the first page of the wizard, enter a name for the IP pool (maximum 100 characters).
3. In the Datacenter tree, select the datacenter where this IP pool will be active.
4. On the Networks page, specify one or more networks for which the pool will be available. Select a network from the Available Networks list and use the down arrow to move it to the Configured Networks list. Each network can be assigned to one IP pool for each datacenter.
5. On the Network Properties page, configure the networking details for the pool. In addition to assigning the IP address, these values will also be configured on the deployed VM (Subnet Mask, Gateway, DNS Suffix, Primary DNS, Secondary DNS).

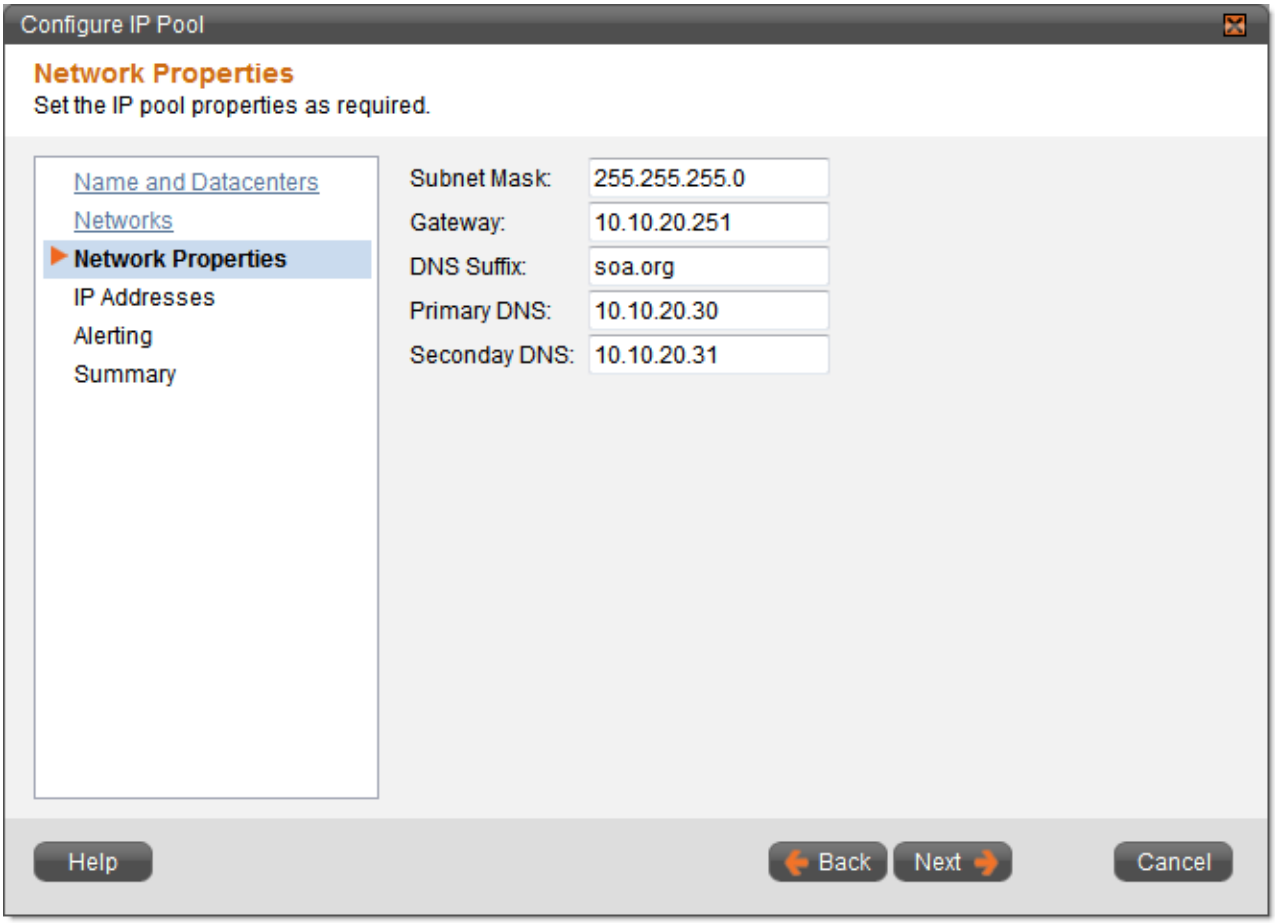

6. On the IP Addresses page, enter the addresses that will comprise the IP pool. Enter single IP address in the From field, and use the From and To fields to define ranges. Click **Add** to add your selections to the pool. IP addresses in a pool do not have to be contiguous.

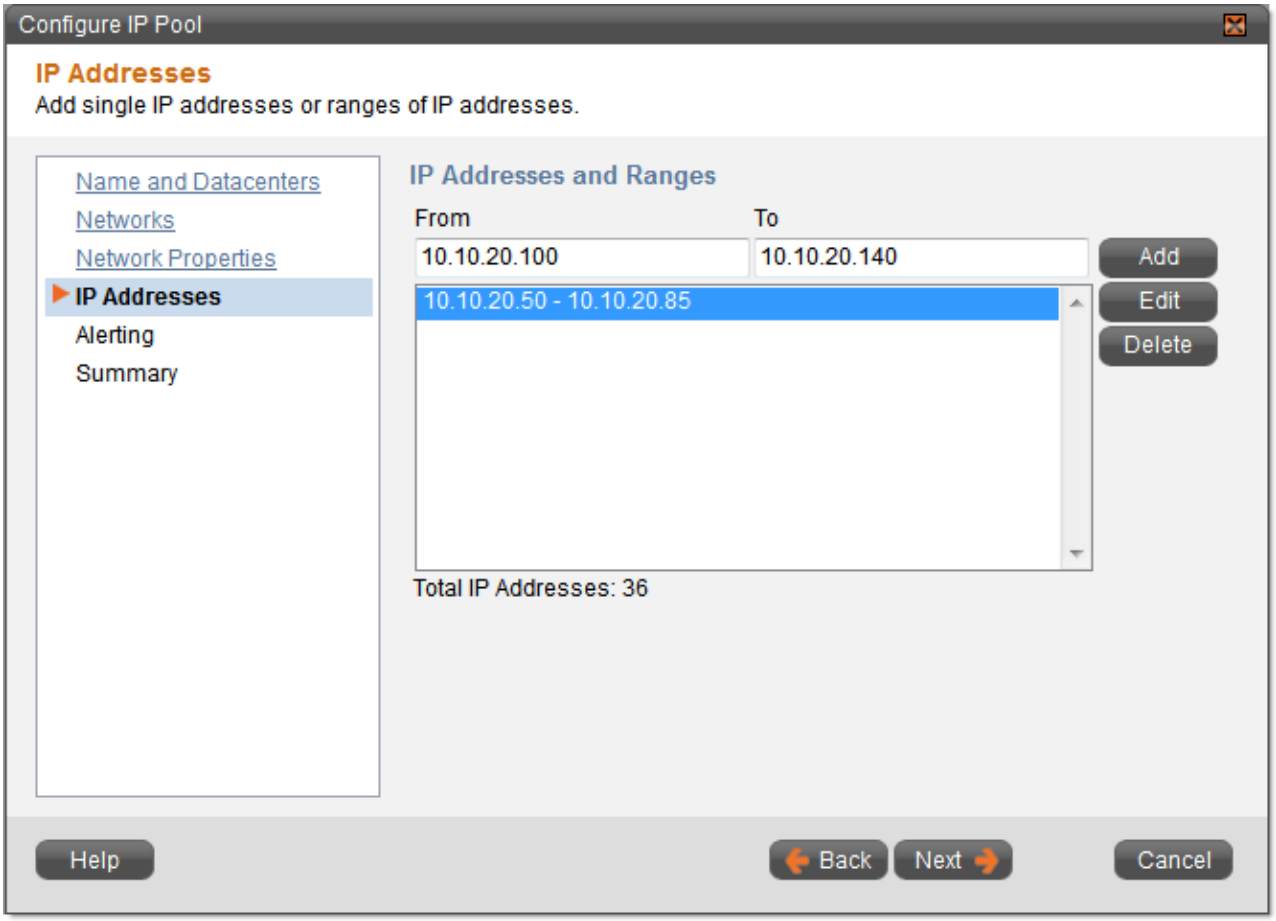

- To avoid IP address conflicts, create the IP address pools so that they cover all the networks that may be routed to each other.
- vCommander does not validate whether a route exists between your selected IP addresses and the gateway, so you must make sure a valid route exists.
- An IP address pool only considers IP addresses in use on the networks that are assigned to the IP address pool.
- If an IP address is used by a VM on a network outside the pool, no warning of a conflict appears for the pool.

To edit the IP addresses, select the individual IP address or the range, click **Edit**, make the required changes, and click **OK**.

7. On the Alerting page, enter the Free IP Threshold. This number determines when vCommander will trigger notifications about the number of free IP addresses remaining in the pool. (Notifications are covered later in this topic.)
8. On the Summary page, review your choices and click **Finish** when satisfied.

Create customization specifications

Customization specifications are XML files used by VMware that contain guest operating system settings for VMs. You create customization specifications in vSphere with the Guest OS Customization wizard, and manage specifications using the Customization Specification Manager in vSphere.

vCommander supports unattended VM customization for all Windows operating systems (using the answer file unattend.xml).

For Windows Server 2003, Windows XP and Windows 2000, vCommander also supports use of the answer file sysprep.inf.

vCommander supports VM customization for any of the Linux versions supported by vCenter.

To assign the IP addresses you've added to the pool, you must deploy the VMs using a customization specification that already exists. vCommander reads the customization specifications available and presents them to you for use when creating services or manually deploying VMs.

See the VMware documentation to learn how to create customization specifications.

Assign customization specifications to your service catalog entries

Access through:	**Configuration menu > Service Request Configuration > Service Catalog tab**
Available to:	**vCommander Role of Superuser and Enterprise Admin**

You can configure the services in your catalog to use addresses taken from specified IP pools. Because you can also control what services are visible to specified users or groups, this provides a perfect opportunity to segregate VMs between groups.

For example, if you have integrated with Active Directory and have groups for your development and QA teams, you can set up IP pools for each team and assign the pools to the services available only to each group.

To update an existing service to use a customization spec:

1. Go to **Configuration > Service Request Configuration > Service Catalog** tab.
2. Click **Edit** under the service to which you want to apply a customization specification.
3. In the Edit Service wizard, click **Next** until you get to the Component page for a guest OS that supports customization.
4. Choose a specification from the Customization Spec drop-down list.

   Troubleshooting: If you know customization specs exist on the vCenter but do not see any on this page, see the Knowledge Base article Why Are Customization Specs Not Appearing in vCommander™?
5. Click **Finish**.

Update the automated deployment placement

Access through:	**Configuration menu > Service Request Configuration > Provisioning Configuration tab > Automatic Deployment Placement pane**
Available to:	**vCommander Role of Superuser or Enterprise Administrator Administrator and All Operator Levels of Access Rights**

You can configure the auto-deployment placements in your provisioning configuration to use specified IP pools. As with controlling services visible to specified users or groups, this provides a perfect opportunity to segregate VMs between groups, because automated placement can also be assigned to users or groups.

To update an existing automated deployment placement to use an IP pool:

1. Go to **Configuration menu > Service Request Configuration > Provisioning Configuration** tab.
2. Select a deployment destination and click **Edit**.
3. In the Edit Automated Deployment Placement wizard, click **Next** until you get to the Network page.
4. Choose between **Connect to the same network as the source service** and **Connect to specified network**. If you choose to specify a network, make sure you also chose it on the Network page of the Configure IP Pool wizard.
5. For **If there is an IP pool linked to target network**, choose **Assign each NIC a static IP address from the pool**.
6. Click **Next** until you reach the Summary page and click **Finish**.

Configure email notification for IP pools

Access through:	**Configuration menu > IP Pools and Networking > Notification tab**
Available to:	**vCommander Roles of Superuser and Enterprise Admin**

vCommander sends email to accounts you specify when the following notification events occur regarding your IP pools:

- **IP Pool Almost Full:** Triggered when the free IP threshold configured for the pool has been reached.
- **IP Pool Full:** Triggered when there are no remaining free IP addresses in the pool.
- **IP Conflict:** Triggered when an IP conflict exists on the network for one of the pool's addresses.

To configure IP pool notification:

1. On the Notification tab, click **Add**.
2. Enter the user or group name and click  .
3. The account details appear. These account details reflect the information that was entered for the user account in **Configuration > Users and Roles**.

   Repeat this step until all required accounts have been added.

   To stop email from being sent to a specific user account, click **Delete**.

4. Click **OK**.

Using IP pools for manual deployment

Access through:	**Views menu > VMs and Templates or Service Requests > Request Details**
Available to:	**Administrator, Operator, Operator with Approval Access Rights**

You can also use IP Pools when manually deploying VMs by using the context menu commands **Provisioning > Clone VM** and **Provisioning > Create Linked Clone**. The same options are presented in the wizards for both commands with respect to customization.

To use IP pools during manual deployment, configure the following settings:

1. On the Resources page, enable **Assign a static IP from pool when an IP pool is linked to the network**.
2. On the Customization page, choose one of the two options which support IP pools:
   - **Customize using wizard**

     With this option, vCommander generates a customization specification that is passed to vCenter at the time of deployment and is then discarded.

     The first page of the customization wizard asks for Domain/Workgroup membership, a Domain Admin's credentials used to join the Domain, Organization, Host and Full (user) Names for the VM, and the Windows Product Key to use. You can also choose to change the SID.

     On the next page you choose the license type and counts, the time zone, administrator password, auto-login and commands to run at system startup.
   - **Customize using a customization specification**

     With this option, you select a customization specification that already exists in vCenter, just as you would when creating or editing a service.

Deleting an IP pool

Access through:	**Configuration menu > IP Pools and Networking > IP Pools tab**
Available to:	**vCommander Roles of Superuser and Enterprise Administrator**

To delete an IP pool, select an IP pool on the list in the IP Pools tab. Click **Delete** and on the Confirm Delete IP Pool dialog, click **Yes**. *Only the IP pool is deleted,* not the IP addresses. The IP addresses are now free to be assigned to another pool.

Managing IP address usage within a pool

Access through:	**Configuration menu > IP Pools and Networking > IP Pools tab**
Available to:	**vCommander Roles of Superuser and Enterprise Admin**

For each IP pool, you can determine:

- what the IP addresses are, including IP address ranges

- whether IP address conflicts exist
- what IP addresses are assigned, reserved, or preempted

See the table below for a list of IP address states.

To view and manage IP address usage:

1. On the IP Pools tab, select the IP pool in the list.

   The IP Pool Details pane lists the main information for the selected IP pool.

2. To display more details, click **More Details**.

   The IP Pool Details dialog lists all IP addresses within the IP pool and the aggregate numbers for the IP addresses.

   Troubleshooting: If the counts are not accurate, a **Repair Counts** button is displayed, as shown in the following image. Click **Repair Counts** and confirm the repair to display accurate counts.

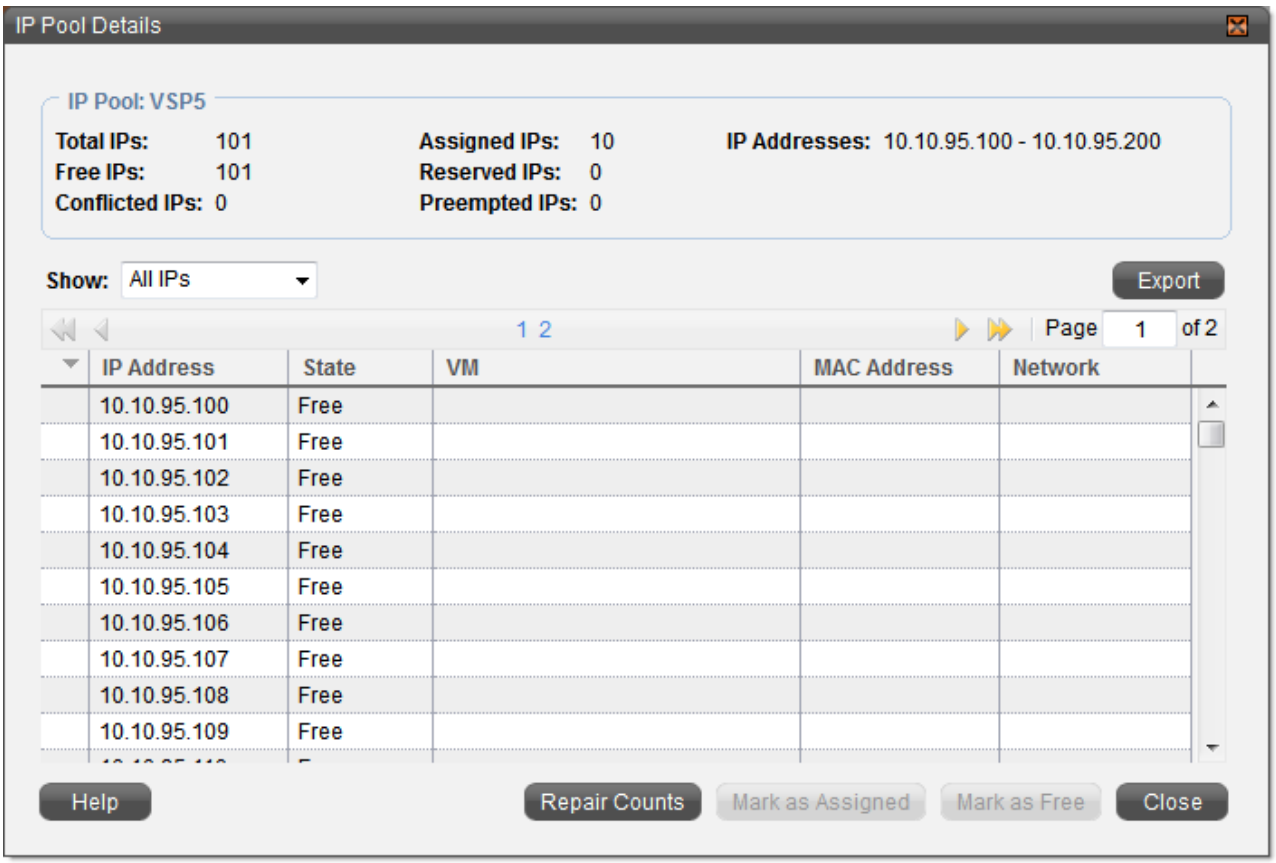

3. To filter the list of IP addresses that are displayed, select a filter from the **Show** drop-down menu.
4. To export the list to CSV format, click **Export**.
5. To mark an IP address that has been assigned to a VM through a manual provisioning process, select it in the list, click **Mark as Assigned** and confirm the change.
6. To mark an IP address as free and available for use, select it in the list, click **Mark as Free** and confirm the change.

7. Click **Close**.

IP Address states within a pool

This state for an IP address:	Means this:
Reserved	The IP address is reserved as a VM is being cloned but has not yet been assigned to the VM. If the clone fails, the IP address is released and becomes available.
Assigned	The IP address has been assigned to a VM.
Preempted	A VM is using an IP address that has not been assigned by vCommander.
Free	The IP address is not being used. Note that when a VM or a NIC is deleted, the IP address becomes free.

What happens when a datacenter or IP pool network is deleted or renamed?

When a datacenter or IP pool network is deleted or renamed, a warning icon () appears beside the IP pool name in the list of IP pools and on the IP Pool Details pane.

Network Fencing

Overview

Embotics® vCommander® provides the ability to deploy vCenter services to isolated networks where IP and MAC address conflicts are not possible, called fenced networks. To do so, a virtual router is deployed with the VMs or virtual services when you define a service as fenced. The deployment of the vRouter is handled entirely by vCommander, so you do not have to import the template or take any other action outside vCommander to make it available for use.

The cost of the vRouter is not presented to users requesting a service which requires one, as the vRouter is an operational cost associated with offering the service. If you need to apply a cost for fencing that users can see, or a cost that will be applied to their quota, use a custom attribute assigned an appropriate cost on the fenced service.

The vRouter is named automatically by vCommander using the following convention:

```
vRouter_R[Request Number]_[Switch Type: (D|S)]_V[vLAN Number]
```

so that you can tell from its name that a vRouter called vRouter_R18_D_V205 was created with service request 18 and uses a distributed switch with vLAN 205. Summary details for each fenced VM optionally display the name of the vRouter isolating it from the rest of your network.

When viewing fenced services in vCommander, the terms external and internal are used to refer to the IP addresses of fenced VMs. The external address is the IP address used to contact the VM from outside the fence, and is known to the router and not the VM itself. The internal address is the IP address actually assigned to the VM's interface, accessible by other VMs inside the fence. When viewing fenced services in the Service Portal, the term Public is used instead of external IP address and Private is used instead of internal.

The external address of the router is assigned by an IP Pool you have created for use with fencing or by a DHCP server available on the external network. This is not user-configurable, so Service Portal users are not required to have any knowledge of the networking involved in order to request and have functional fenced networks provisioned transparently.

The private IP addresses for VMs inside the fence can be assigned statically, dynamically, or both. Each network interface for each component of the service is independently configurable.

The public addresses for the components inside a fence all come from IP pools, even if the router gets its address from DHCP.

To assign the IPs statically, you provide the IP address the service will use. When dynamic addressing can be used, the VMs will receive their addresses from a DHCP server on the router, which means the service's templates must be DHCP ready.

Each service that employs network fencing requires a dedicated vLAN. Both standard and distributed switches are compatible with network fencing.

You can further configure the fence to allow communications in multiple supported configurations. Each network interface for each component of the service is independently configurable with an access mode that defines the default firewall rules that apply. The access modes are:

- **Out only:** No inbound connections are permitted on any port, but outbound connections are allowed on all ports.
- **In and Out:** All inbound and outbound connections are permitted on all ports.
- **In only:** All inbound connections are permitted on all ports, but no outbound connections are allowed on any port.
- **None:** No inbound or outbound connections are permitted on any port.

When you choose an access mode of **None** or **Out Only**:

- You can open specific ports for inbound connections as exceptions to the default rule applied by the selected access mode. By default, traffic for a fenced VM with In or In/Out Access is redirected to all of its private ports.
- All access to the VMs must be handled via console connections, unless you have opened specific ports for inbound connections. When inbound communications are required, the network must be configured for promiscuous mode.

Important: Networking fencing is only available for services deployed to VMware vCenter 5.0 or later managed systems.

Here are the major steps you carry out to set up network fencing:

1. Prepare vCenter back-end networking. This involves setting up non-routable VLANs for the fenced networks.
2. Create an IP pool dedicated to fencing.
3. Create a deployment destination for fenced networks.
4. Create a service using a fenced network.

Adding and removing DNS records for fenced VMs

When you deploy a fenced service, the VMs behind the fence are segregated from the rest of your network. This means that DNS records are not automatically created for fenced VMs. Users outside the fence are still able to connect, providing you've allowed IN access, but only by IP address.

Decommissioning fenced networks

No special consideration is required to decommission the vRouters handling a fenced network. When all VMs inside the fence have been removed, vCommander automatically decommissions the vRouter.

Best Practice: Deleting the fenced service at the virtual service level is recommended.

Connecting to the vRouter

If you need to connect to the deployed vRouter in a fenced network, use console access. SSH access has been disabled for security reasons.

Preparing vCenter Networking for Fencing

Preparing networking is the first step in configuring fenced networks. See Network Fencing for an overview of the entire process.

Each fenced service you deploy requires a dedicated vLAN. Because the vLANs cannot be shared, it’s important to configure your managed system with enough vLANs to ensure that service fulfillment is possible.

Embotics® vCommander® assigns the next available vLAN from those available in the deployment destination.

If your deployment destination is a cluster, you must either have a distributed switch configured for the cluster, or each host in the cluster must have an identical VLAN portgroup configured on each host with a standard switch.

Best Practice: Create a block of non-routable vLANs with no network services available on the backing switch or switches that will be used only for fencing. During deployment, vCommander will create the portgroups on the desired vLAN based on the configuration of the placement destination.

When configuring the network, if you're using standard switches, make sure to set the security policy on the vSwitch to accept promiscuous mode so that it is ready to handle all service requests, no matter what communications configuration has been selected. If you use distributed switches, vCommander creates the portgroup and configures it appropriately.

What's next?

Continue to Creating an IP Pool Dedicated to Fencing.

Creating an IP Pool Dedicated to Fencing

Access through:	**Configuration menu > IP Pools and Networking > IP Pools tab**
Available to:	**vCommander Roles of Superuser and Enterprise Administrator**

Creating an IP pool is the second step in configuring fenced networks. See Network Fencing for an overview of the entire process.

The external IP addresses for each VM inside the fence come from an IP pool you configure for this exclusive purpose. External IP addresses are those used to access the fenced VMs from outside the fenced network. The external address of the router is set either from the same IP Pool, or from a DHCP server on the external network, depending on the configuration of the service.

Embotics® recommends choosing a name for the IP pool that will allow you to easily identify it as designed for use with fencing.

Follow the steps below to create an IP pool.

1. Go to Configuration > **IP Pools and Networking** and select the **IP Pools** tab.
2. Click **Add**.
3. On the first page of the wizard, enter a name for the IP pool and choose the datacenter where it will be active.
4. On the Networks page, check all of the networks for which the pool will be available. Each network can be assigned to one IP pool for each datacenter.
5. On the Network Properties page, configure the networking details for the pool. In addition to assigning the IP address, these values will also be configured on deployed VMs (Subnet Mask, Gateway, DNS Suffix, Primary DNS, Secondary DNS).
6. On the IP Addresses page, enter the addresses that will comprise the IP pool. Enter a single IP address in the From field; use the From and To fields to define ranges. Click **Add** to add your choices to the pool. IP addresses in a pool do not have to be contiguous.
7. On the Alerting page, enter the Free IP Threshold. This number determines when Embotics® vCommander® will trigger notifications about the amount of remaining free IP addresses in the pool.
8. On the Summary page, review your choices and click **Finish** when satisfied.

Best Practice: You can create more than one IP pool to use with fencing, but Embotics® recommends that you complete the steps required to get network fencing functional with a single IP pool, deployment destination, and service before introducing more complexity to your configuration.

What's next?

Continue to Creating a Deployment Destination for Fenced Networks.

For more information on using IP pools, see Configuring and Managing IP Pools.

Creating a Deployment Destination for Fenced Networks

Access through:	**Configuration menu > Service Request Configuration > Provisioning Configuration tab > Automatic Deployment Placement pane**
Available to:	**vCommander Role of Superuser or Enterprise Administrator Administrator and All Operator Levels of Access Rights**

Creating a deployment destination is the third step in configuring fenced networks. See Network Fencing for an overview of the entire process.

Deployment destinations used with fencing must provide network capabilities to each VM deployed as part of the fenced service. This means that when a cluster is configured as a destination for fenced services, you must either have a distributed switch configured for the cluster, or each host in the cluster must have an identical VLAN portgroup configured on each host with a standard switch. Embotics® recommends choosing a name for the deployment destination that will allow you to easily identify it as designed for use with fencing.

When the virtual router is deployed for a service, no ownership is assigned as part of the request fulfillment. This means Service Portal users will not be able to see the router, typically the desired behavior. However, ownership can be assigned to the vRouter by a Embotics® vCommander® user with an appropriate role.

Existing deployment destinations may also be reconfigured to use fencing by choosing to edit them and following the same steps. The fencing configuration for deployment destinations will only be used if the published service they are deploying is configured with fencing.

1. Under the Configuration menu, choose **Service Request Configuration**. Switch to the **Provisioning Configuration** tab.
2. On the Automated Deployment Placement pane, click **Add**.
3. On the Name page, enter a name that identifies the destination.
4. On the Managed System page, select a vCenter managed system from the list.
5. On the Assignment page, you set up the destination by user account:
   - to create a destination for all users, select **Default Destination**;
   - to create a destination for one or more specific user accounts, switch to the **Users/Groups** tab and enter the login/email account information, then click **Add**;

- to create a destination for one or more specific organizations, select an organization from the drop down menu and click **Add**.

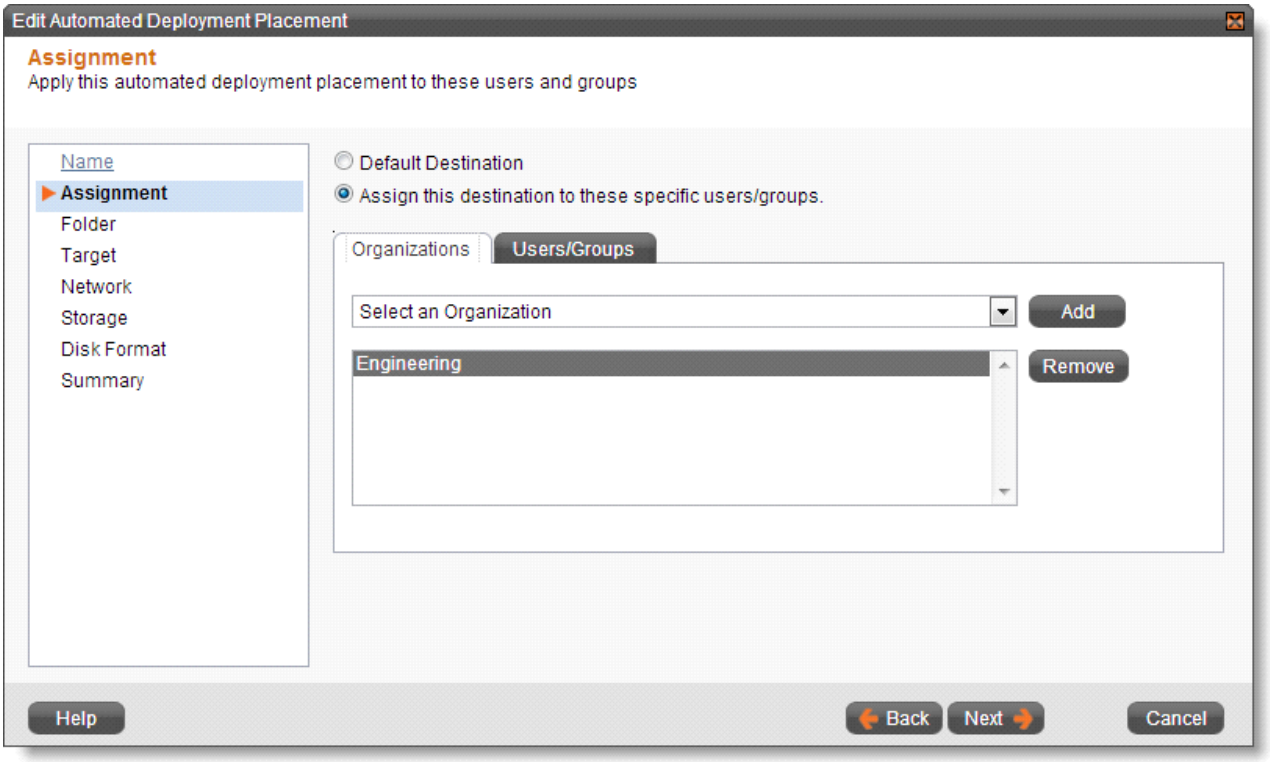

You can add both users and groups as well as one or more organizations to a single destination.

6. On the Folder page, select the appropriate folder in the tree.
7. On the Target page, select the appropriate target host or cluster in the tree. If you have chosen a cluster that has its automation level for VMware DRS (load balancing) set to Manual, you can select a host within that cluster. Otherwise, this selection is unavailable.
8. Choose to deploy based on Peak Capacity or Average Capacity by selecting from the drop-down menu. For more information, see Managing Host and Cluster Capacity.
9. On the Networks page, you define the network to which VMs will attach once deployed.

   By default, **Connect to the same network as the source service** is disabled. Enable this option if you want to match the source template's settings.

   Otherwise, select one or more networks from the list of available networks and click the down arrow (or double-click) to move them to the Configured Networks list.

   See Configuring Networks to learn how to assign networks to zones. Note that when you allow users to choose the network zone on the request form, you must add at least one network from each zone to the deployment destination, or automated deployment will fail.

   The $NETWORK deployment parameter specified by a request approver or in a script overrides the network zone specified on the request form.
10. On the IP Pools page, choose whether or not to assign IP addresses from a pool. These settings are not used by fenced services deployed to this destination, but are provided to allow fenced and non-fenced services to share a deployment destination.

If an IP pool is linked to a target network:

- To have vCommander automatically assign a static IP address to each NIC, select **Assign each NIC a static IP address from pool**. For this option, a customization specification must have been configured for the source service. For more information, see Manually Provisioning vCenter Services and Managing the Service Catalog.
- To allow IP addresses to be assigned manually during deployment, select **Do not assign any IP addresses**.

11. On the Isolated (fenced) Networks page, select **Enable fenced networks** to allow fenced deployments to be placed on this destination, and make a selection from the drop-down menu.

    If you have a large number of networks, click next to the drop-down menu to open a pop-up dialog for easier searching. Mouseover the question mark to learn what properties are searched.

    On this page you define the physical network bindings with which fenced VMs will be configured when deployed. Choose either **Distributed Switch** or **Standard Switch** to match your environment.

    - If your environment is configured to use a distributed switch, select it from the list of available switches.
    - If your environment requires a standard switch, select the virtual switch for each host. Most typically, you will set the same vswitch on each host.

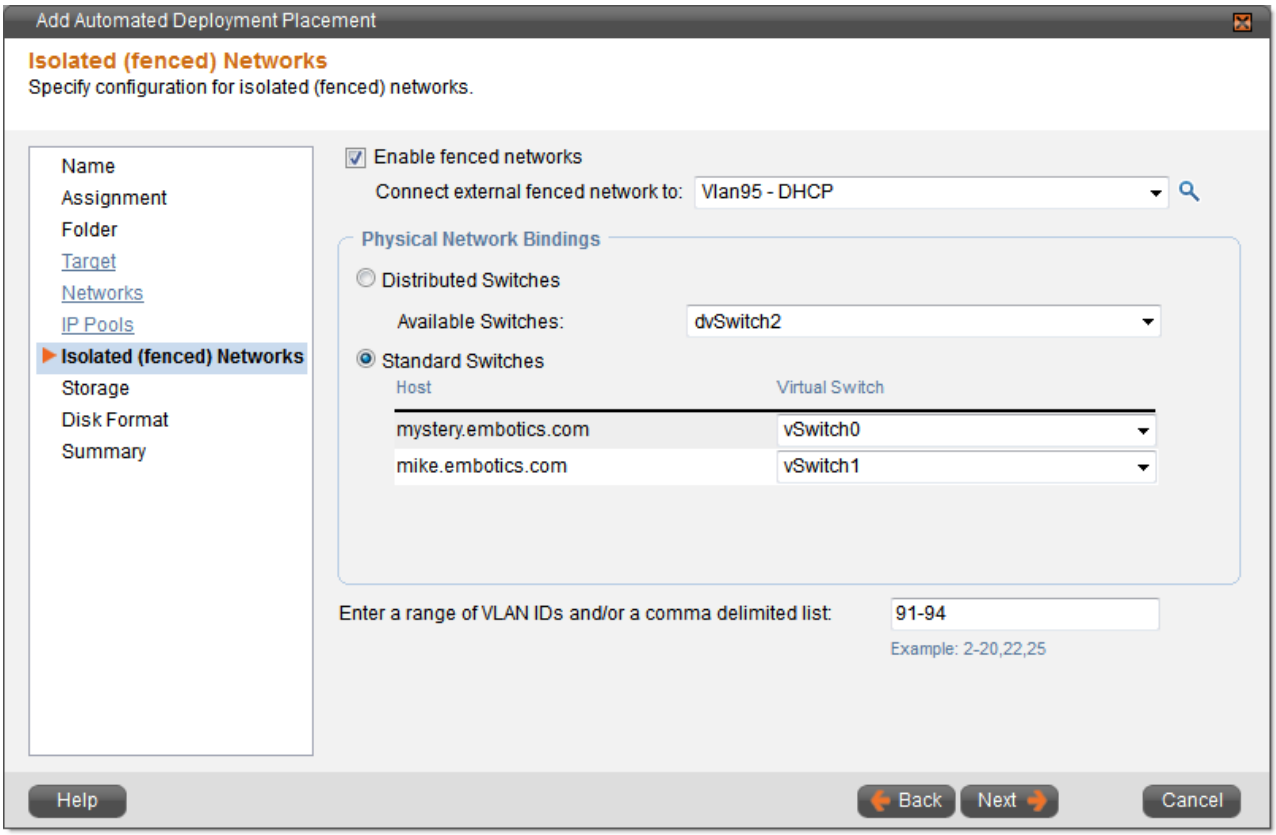

Lastly, enter IDs for the vLANs that have been configured on the host uplinks and reserved for fencing by yourself or your system administrator. Enter a comma-separated list of one or more vLAN IDs and/or a range of vLAN IDs available to the destination cluster. If you aren't sure what vLANs are available, contact your network administrator or consult the vendor documentation for how to retrieve the information by logging into the switch.

12. On the Storage page, select a datastore from the list of available datastores and click the down arrow to move it to the Configured Datastores list. If required, select other datastores to add to the list of configured datastores.

    Making multiple datastores available:

    - ensures that VMs can be automatically deployed to datastores that have free space available
    - allows datastores with different storage tiers to be selected for automated deployment
    - ensures that VMs with disks on different storage tiers can be deployed to meet service level agreements

    Datastores that have been excluded from capacity calculations are not displayed on the list of available datastores. For more information, see Managing Host and Cluster Capacity.

13. On the Disk Format page, select the format in which to store the VM's virtual disks.
14. On the Summary page, review your options and click **Finish**.

For standalone hosts only, the host property appears as "Not set" on the Summary page. Any standalone host will also appear as "Not set" on the Details pane for the placement destination you created.

What's next?

Continue to Creating a Service Using a Fenced Network.

Creating a Service Using a Fenced Network

Access through:	**Configuration menu > Service Request Configuration > Service Catalog tab**
Available to:	**vCommander Role of Superuser, Enterprise Admin**

Creating a fenced service is the final step in configuring fenced networks. See Network Fencing for an overview of the entire process.

With the IP pool and deployment destinations configured for network fencing, you can now create services ready for isolation. Existing services may also be reconfigured to use fencing by choosing to edit them and following the same steps.

1. Under the Configuration menu, choose **Service Request Configuration**. Switch to the **Service Catalog** tab.
2. Click **Add Service**.
3. On the Service Description page, provide a name and description for the service. Including details to indicate that the service is fenced ensures Service Portal users have the information they need to make appropriate decisions. A custom icon provides a visual cue that the service is fenced, and a category for fenced services is also recommended.
4. On the Components page, click **Add** and choose either **VMs and Templates** or **Virtual Services**.

5. Enable **Allow VMs in the catalog** if the source of the service you want to add is not a template. Use the VMs and Templates view and/or the search to refine the list of VMs or virtual services displayed, select your choice from the list, and click **Add to Service**. Once you have added all required components, click **Close**.
6. Enable **Deploy as fenced** to have the service deployed in an isolated network.
7. A subpage is added to the Add Service wizard for each component you added to the service in step 6. On each subpage, you see a button bar containing at least **Infrastructure**, **Resources**, **Attributes** and **Form** buttons. These buttons allow you to customize options for each component in the service. Adding elements to the Form tab allows requesters to change the default settings you configure on the other tabs. For more information on component-specific settings, see Adding a vCenter Service to the Catalog.
8. On the Deployment page, choose to deploy the service as individual components or as a virtual service. If you choose to deploy the service as a virtual service, you must also choose whether to use the default naming format or override it with one you define.
9. Use the arrow controls to choose the deployment order of the components. When deployed as a virtual service, this order will also be the startup order used when powering the service on. Powering the service off will use the reverse order.
10. Select a completion workflow for the entire service, if required.
11. When you select to deploy as a fenced network on the Components page, the Fencing Configuration page appears. On this page, you define how the networking for the components will be defined on the vRouter.
    - In the External Router Interface pane, the **IP Address Assignment** setting controls how the public external addressing for components is assigned. Choose **DHCP** to dynamically allocate addresses from a DHCP server on the network to which you're deploying the fenced service, or choose **From IP pool** to have vCommander allocate an address from the IP pool configured for the deployment destination.
    - In the Internal Router Interface pane, use the **IP Address** field to enter the address to assign as the gateway for the components on the vRouter deployed with the service.
    - Use the **Subnet Mask** field to configure a subnet mask other than the default 255.255.255.0.

    For each component, the NICs are identified in the Component Interface Configuration list.
12. Set the Access mode for each NIC:

    In/Out - Communication to and from the external public network is allowed on all ports.

    In Only - Inbound communications from the external public network are allowed on all ports, but no outbound communications are allowed on any port.

    Out Only - Outbound connections to the external public network are allowed on all ports, but no inbound communications are allowed on any port. Click **Add Ports** to create port forwarding rules for specific exceptions. See below for more information.

    None – No communications to or from the external public network is allowed on any port. Click **Add Ports** to create port forwarding rules for specific exceptions. See below for more information.
13. Set the Address Mode for each NIC, which defines how fenced components obtain IP addresses.

    DCHP – The IP address will be assigned by the router providing the fence.

Static – The address you provide in the Static IP Address field will be used. vCommander assigns the specified IP address to that component if requested over DHCP. The base image can be configured either with this static IP address, or with DHCP.

Static Assign – The address you enter in the Static IP Address field will be used. Entering the address provides the router with information needed to create a route to the VM, and vCommander with the information to use during customization. If a customization spec is not assigned to component in the service catalog, vCommander will automatically create one and apply the static IP address through it. You can optionally add a Host Name; otherwise, the VM name will be used in the guest OS.

14. By default, the vRouter in a fenced network redirects traffic from all public ports on the vRouter to all private ports on the fenced VM. You may want to expose only specific ports; specific application services in the fence can then be contacted via the exposed ports. For example, if one of the fenced VMs is a web server, you might want to expose only ports 80 and 443 so that clients on the external public network can access the website, while still blocking all other inbound communications.

    To expose only specific ports, under Port Forwarding, click **Add Ports**.

    To configure port forwarding for a component, its Access Mode must be **Out Only** or **None**.

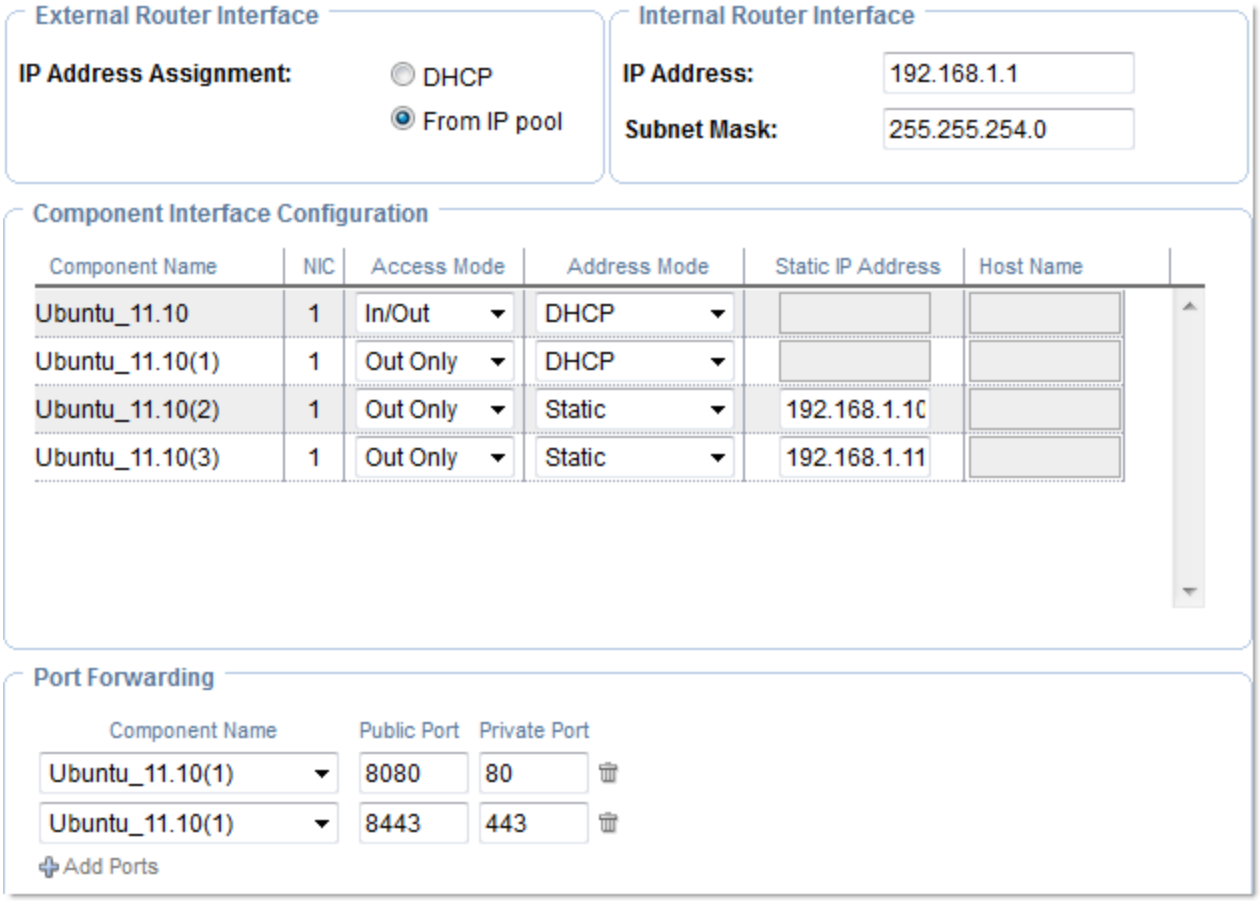

If your service has multiple components, in the Component Name drop-down menu, select the component you want to configure.

In the Public Port field, enter a port number. The public port is the port on the vRouter that external traffic will connect through.

In the Private Port field, enter a port number. The private port is the port on the fenced VM that the vRouter will direct traffic to.

You can specify a public port only once, but multiple public ports can forward to the same private port. For example, you can forward public ports 80 and 443 to private port 443.

See Viewing Fencing Information for VMs to learn how users can view the public IP address for accessing the VM when port forwarding is configured.

15. On the Visibility page, specify who is able to request the service:
    - Selecting **Do not publish** allows you to complete the configuration of the service without making it available for anyone to request. You can also use this selection to render a service temporarily unavailable by editing it later.
    - Selecting **Publish – Global** makes the service available to all users allowed to request services.
    - Selecting **Publish – Specific organizations, users and groups** allows you to limit who can see the service to those users you define.
16. On the Summary page, review your choices and click **Finish**.

Viewing Fencing Information for VMs

This article shows you how to view information necessary for connecting to fenced VMs.

See Network Fencing for an overview of configuring fenced networks.

Fenced VM - General pane

To connect to a fenced VM, users must know the VM's public IP address.

In both Embotics® vCommander® and the Service Portal, users can go to the General pane on the VM's Summary tab. The **IP Address** property shows the VM's public IP address.

When the access mode is Out Only or None, the IP address field shows the IP address of the vRouter. If inbound communications are allowed, each VM has its own public IP address.

When you have configured port forwarding for a fenced VM, the IP Address property displays a **Port Forwarding** link in both vCommander and the Service Portal.

Click **Port Forwarding** to display details. The **Public IP** column shows the VM's public IP address. In this example, public port 8080 has been forwarded to private port 80 on the VM, and public port 8443 on the vRouter has been forwarded to private port 443 on the VM. Users must access the VM with the public IP address and public port(s).

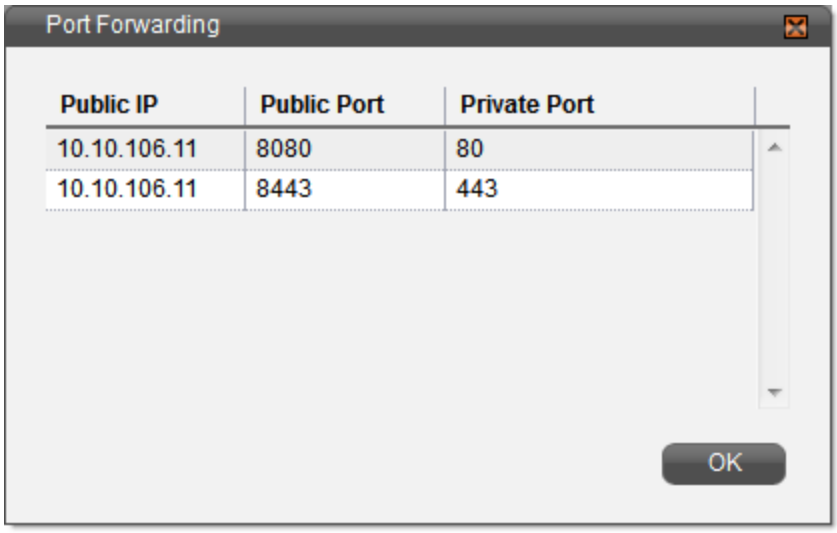

Fenced VM - Details pane

The Details pane of the VM's Summary tab displays network details.

Best Practice: To provide as much at-a-glance information about fencing on the VM's Summary tab as possible, add the following properties to the Details pane by clicking the gear icon in the upper right corner of the pane:

- Fence Name
- Fenced - whether the VM is in a fenced network

- vRouter - whether the VM is a vRouter for the fenced network

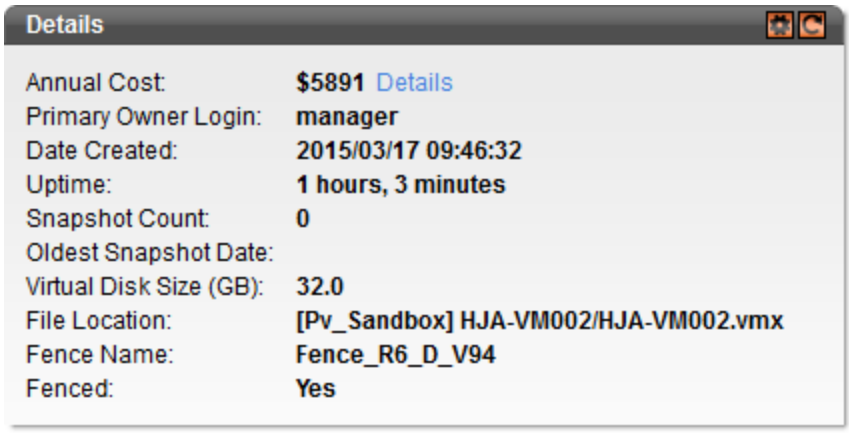

vRouter - General pane

The Summary tab for the vRouter VM also provides useful details on the General pane, such as its public IP address, private IP address and vLAN number:

Integrating BlueCat™ IP Address Management with vCommander

You can configure Embotics® vCommander® to use BlueCat™ IP address management (IPAM) capabilities. When you integrate vCommander with BlueCat™ IPAM as shown in this article, you can create host records or DHCP reservations from BlueCat™ IPAM during provisioning with vCommander.

To set this up, you need to:

- Copy the BlueCat™ API bindings file to the vCommander installation directory.
- Create a set of credentials that has access to the BlueCat™ server.
- Connect vCommander to a BlueCat™ server.

- Set up a completion workflow with a Configure OS Networking step, or add this step to an existing completion workflow.

✏ The Configure OS Networking step is also supported in command workflows.

- Set up a decommissioning workflow to release the addresses when decommissioning VMs managed by BlueCat™.

These tasks are explained in this topic. For an end-to-end example that includes BlueCat™ networking, see Automating VM Customization through Workflows: Examples.

Copy the BlueCat™ JAR file to the vCommander installation directory

1. Download the BlueCat™ API bindings file (likely called api.tgz) from the BlueCat™ website.

   If you cannot locate this file, contact your BlueCat™ representative.
2. Create a subdirectory named "ext" in the following directory:

   `<vCommander_install_dir>\tomcat\lib`
3. From the downloaded file, extract the `api.jar` file to the following directory:

   `<vCommander_install_dir>\tomcat\lib\ext`
4. Restart the vCommander Windows service.

Create system credentials for the BlueCat™ server

Access through:	Configuration menu > Credentials
Available to:	vCommander Role of Superuser and Enterprise Admin

1. On the Credentials page, click **Add**.
2. Enter the user name and password for a BlueCat™ account.

   The vCommander integration with BlueCat™ automatically uses `sudo` with non-root-user credentials, so the user must be able to run `sudo` in this case. vCommander supports interactive `sudo` prompts. Ensure that the `sudoers` file on the BlueCat™ server does not contain the following line: `Defaults requiretty`
3. For the description, enter "BlueCat", to serve as a memory aid to administrators when configuring tasks requiring credentials.
4. In the Category drop-down list, select **System Credentials**.
5. Click **OK**.

Connect vCommander to a BlueCat™ server

Access through:	Configuration menu > System Configuration > Integration tab
Available to:	vCommander Role of Superuser

1. On the Integration page, click **Add > BlueCat™ IPAM Server**.
2. Enter the host name or IP address and port for the BlueCat™ server.

   Both http and https protocols are supported. The default port is 80.

3. Select the BlueCat™ credentials in the **Credentials** drop-down menu.
4. In the **Configuration** field, enter the BlueCat™ configuration (the parent IPAM object).
5. In the **Server** field, enter the DNS/DHCP server to deploy the changes to. You can add multiple servers as a comma-separated list.
6. For **Services**, select **DNS** and/or **DHCP** as required.
7. Click **Test** to test the connection.

   vCommander queries BlueCat™ to determine whether everything is set up correctly. If you see an error about a missing file, see Copy the BlueCatTM API Bindings to the vCommander Installation Directory above.

   vCommander connects to BlueCat™ only as needed; there's no continuous connection.
4. Once you see a Success message, click **OK** to save the configuration.

Set up a completion workflow

Access through:	**Configuration menu > Service Request Configuration > Completion Workflow tab**
Available to:	**vCommander Role of Superuser and Enterprise Administrator**

1. Create a completion workflow.
2. On the Name page, provide a name, and in the Apply this workflow drop-down menu, select **after a VM is deployed**.
3. On the Steps page, add a **Configure OS Networking** step.
4. Give the step a descriptive name.
5. In the **Credentials** drop-down menu, select credentials or click **Add Credentials**. See Configuring OS Networking through a Workflow Step for guidance.
6. In the **Assign IP** drop-down menu, select **From BlueCat™ IPAM**.
7. Enter networking details as required. The Gateway, DNS View, Block, Network and Domain Name fields are mandatory.
8. In the Action drop-down list, select **Apply Settings** to configure the VM in this step, or select **Reserve Only** to have the IP information recorded in the approval comments. When you select **Reserve Only**, the output can be used as input to a subsequent Customize VM step in the workflow.

Workflow errors are written to the workflow step comment log. See Troubleshooting and Tracking Workflows for more information.

For full details on completion workflows, see Creating a Completion Workflow.

Set up a decommissioning workflow

Access through:	**Configuration menu > Service Request Configuration > Completion Workflow tab**
Available to:	**vCommander Role of Superuser and Enterprise Administrator**

When you use vCommander to create host records or DHCP reservations from BlueCat™ IPAM during provisioning, it makes sense to release the addresses when decommissioning VMs.

1. Create a new completion workflow.
2. On the Name page, provide a name, and in the Apply this workflow drop-down menu, select **after a Change Request is fulfilled**.
3. On the Steps page, add a Decommissioning Networking step (in the Guest OS category).
4. On the Assigned Forms page, select **Apply this workflow to the selected forms** and select the Decommissioning Request form.

For full details on completion workflows, see Creating a Completion Workflow.

Disabling or removing the BlueCat™ server

Access through:	**Configuration menu > System Configuration > Integration tab**
Available to:	**vCommander Role of Superuser**

Disabling an external server makes the server unavailable for connections but saves the settings, meaning that you can return to the configuration dialog later and simply re-enable it.

Removing an external server clears the settings, meaning that you must reconfigure all of the settings if you want to reintegrate later.

To disable an external server

1. On the Integration page, locate the external server you want to disable and click **Edit**.
2. Deselect the Enabled checkbox and click **OK**.

To remove an external server

1. On the Integration page, locate the external server you want to clear and click **Remove**.
2. Confirm the change by clicking **Yes**.

Service Request Management

Learn how to handle service requests:

- Approve or reject service requests
- Specify deployment parameters when approving a service request
- configure maintenance windows for disruptive change request fulfillment
- Fill and track service requests
- Release a VM or virtual service
- Request a service or a change to a service