



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

# WebRTC Private Edition Guide

[Configure WebRTC](#)

---

## Contents

- [1 Override Helm chart values](#)
- [2 Configure Kubernetes](#)
- [3 Configure security](#)
  - [3.1 Arbitrary UIDs in AKS](#)
- [4 Configure the service](#)

---

Learn how to configure WebRTC.

### Related documentation:

- 
- 
- 

### RSS:

- [For private edition](#)

## Override Helm chart values

Download the WebRTC Helm charts from JFrog using your credentials. Override the configuration parameters in the **values.yaml** file to provide deployment-specific values for certain parameters. You can override values in the Helm charts to configure Private Edition. For more information about overriding Helm chart values, see the "suite-level" documentation about how to override Helm chart values: [Overriding Helm chart values](#)

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
deployment.namespace	Name of Kubernetes namespace for WebRTC deployment	mandatory	webrtc	string	You can modify the default namespace used to deploy applications in the deployment.namespace option.	deployment: namespace: production
deployment.priorityClassName	Name of the priority class for pods that specify the importance of a pod relative to other pods	optional		string		
deployment.nodeSelector	Node selector for CoTurn and CoTurn pods	optional		Specification		deployment: nodeSelector: genesysengage.com/

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
						nodepool: general
deployment.tolerations	Include this parameter in the Gateway and CoTurn, if the content of toleration exists.	optional		Specification		deployment: tolerations: - operator: Exists  effect: NoSchedule key: "k8s.genesysengage.com/ nodepool"
deployment.ingress.domain	Ingress domain	mandatory		string		deployment: ingress:  domain: apps.vce- c0.eps.genesys.com
deployment.ingress.annotations	WebRTC Annotation for Ingress controller	mandatory		Specification	As the default value of the HAProxy route timeout is set to 30 s, there is a possibility it interferes with the WebRTC long-polling timeout (30 s) and disconnect the session.	deployment ingress:  annotations:  kubernetes.io/ ingress.class: nginx01-internal  nginx.ingress.kubernetes. affinity: cookie  nginx.ingress.kubernetes. affinity- mode: persistent  nginx.ingress.kubernetes. ssl- redirect: "false"  nginx.ingress.kubernetes. session- cookie- path: "/; Secure"  nginx.ingress.kubernetes. session-

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
						<pre> cookie- samesite: None {{!}}- {{!}}{{!}}deployment.ingr {{!}}{{!}}If this option is defined, tls option is declared in the Ingress specification {{!}}{{!}}optional {{!}}{{!}} {{!}}{{!}}Specification {{!}}{{!}} {{!}}{{!}}deployment:   ingress:     tls:  secretName: webrtc.api01-eastus2.dev. tls-secret </pre>
deployment.affinity	Pod affinity descriptions	optional		Specification		<pre> deployment:   affinity:  podAntiAffinity:  preferredDuringScheduling - weight: 100  podAffinityTerm:   labelSelector:   matchExpressions: - key: servicename  operator: In  values: - webrtc- gateway - webrtc- coturn </pre>

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
						topologyKey: failure- domain.beta.kubernetes.io/ zone
deployment.dnsPolicy	Kubernetes DNS Policy applied in the Pods	optional				deployment:  nodeSelector:  genesysengage.com/ nodepool: general
deployment.dnsConfig	All DNS settings must be provided using the dnsConfig field in the Pod specification	optional				deployment:  dnsConfig:  options: - name: ndots value: "3"
deployment.keda	Enable KEDA usage for the Gateway and CoTurn horizontal auto-scaling	optional	false	true/false		
deployment.coturnDeployment	Type of CoTurn deployment - internal: the internal LBs are created and the IP addresses of that LBs must be used by the firewall or other ways to be exposed externally. external: the external LBs are created with given external	mandatory		internal/ external	For Premise Edition - This parameter is configured as external	

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	static IPs (IPs for the green and blue LBs must be set with lbIpBlue and lbIpGreen during the infra-color deployment.					
deployment.coturnService.annotations	Annotation that is added to the Kubernetes LoadBalancer Service object	optional				<pre> deployment:   coturnService:     annotations:       service.beta.kubernetes.io/ azure-load-balancer-resource-group: service-webrtc-westus2-dev </pre>
monitoring.enabled	Enable monitoring content - dashboards, alerts, metrics	optional	false	true/false		
monitoring.dashboards	Enable ConfigMaps deployment that contains dashboards	optional	false	true/false		
monitoring.prometheusMetrics	Enables Prometheus Metrics deploy PodMonitors	optional	false	true/false		
monitoring.prometheusAlerts	Enable Prometheus Alerts rules for alerts	optional	false	true/false		
image.imagePullSecrets	Secrets to use for image, list	mandatory				<pre> image:   imagePullSecrets: </pre>

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
						- myRegistrySecret
image.pullPolicy	Kubernetes pull policy of all containers	optional	Always	Always/ IfNotPresent		
image.initContainerImage	Image for initialization container - used to create log folders. If image is not specified, the init container is not applied and the logs are written into logPath	optional		string		
image.webrtc	Repository/ directory to get the Gateway image	mandatory		string		pureengage- docker- staging.jfrog.io/ webrtc
image.coturn	Repository/ directory to get the CoTurn image	mandatory		string		pureengage- docker- staging.jfrog.io/ webrtc
image.webrtcVersion	Versions of the WebRTC Gateway container	mandatory		string		9.0.000.88
image.coturnVersion	Versions of the CoTurn container	mandatory		string		9.0.000.88
gateway.replicas	Number of Gateway pods on the deployment stage	optional	1	integer		
gateway.workersCount	Number of Gateway worker threads that handle calls. 1 worker	optional	3	integer		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	handles 25 registrations/calls. CPU and Memory request depends on the number of workers.					
gateway.voiceSipProxy	Voice microservice - SIP proxy address	mandatory		string, address		voice-sipproxy.voice.svc.cluste
gateway.turnExternalBlue	FQDNs of External Blue LB	mandatory		string, address		
gateway.turnExternalGreen	FQDNs of External Green LB	mandatory		string, address		
gateway.authRedirectUri	GWS/WEE redirect URI for WVE authentication	mandatory		string, address		
gateway.authService	GAuth service address	mandatory		string, address		
gateway.envService	GWS9.x Environment service address	mandatory		string, address		
gateway.cfgService	GWS9.x configuration service address	optional		string, address		
gateway.enableTranscoding	Enable or disable transcoding on the Gateway side. Transcoding is enabled by default. If the transcoding is disabled, the Gateway can handle more agent sessions but OPUS codec is not	optional	true	true/false		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	supported.					
gateway.enableIpcCalls	Specifies if the IpcCalls operations are enabled	optional	false	true/false		
gateway.arguments	Any additional options that are applied to the Gateway containers	optional		Array of strings		<pre>gateway: arguments: [ '-codecs pcmu,pcma,opus=120', '-sip- disallowed- codecs opus,telephone- event' ]</pre>
gateway.podAnnotations	Any additional annotations that are applied to the Gateway pods	optional				<pre>gateway: podAnnotations: prometheus.io/ scrape: "true"  prometheus.io/ port: "10052"  prometheus.io/ path: "/metrics"</pre>
gateway.resources	Describes the resources requested for the Gateway pods. <b>Important</b> Do not specify this option, if you do not need resources requests/limits.	optional		Section		<pre>gateway: resources: requests: cpu: 800 memory: 150 limits: memory: "8Gi"</pre>
gateway.resources.requests.cpu	Requested amount of CPU in milliunits.	optional	800	integer		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	<p><b>Important</b></p> <p>This value is per worker and is multiplied by the <code>gateway.workers</code> option in helm</p>					
<code>gateway.resources.requests.memory</code>	<p>Requested amount of Memory (in MB).</p> <p><b>Important</b></p> <p>This value is per worker and is multiplied by the <code>gateway.workers</code> option in helm</p>	optional	150	integer		
<code>gateway.resources.limits.memory</code>	<p>Absolute value for Gateway memory usage limit</p>	optional	"8Gi"	Kubernetes value for the resource limit		
<code>gateway.scaling</code>	<p>Describes the auto-scaling parameters. If the <code>deployment.keda</code> option is set to false, you can skip this option.</p>	optional		Section		<pre>gateway:   scaling:    pollingInterval:   30    maxReplicaCount:   100    prometheusAddress:   http://monitoring-   prometheus-   prometheus.monitoring:9090    thresholdSignins:   70</pre>
<code>gateway.scaling.prometheusAddress</code>	<p>Describes the auto-scaling parameters. If the <code>deployment.keda</code> option is set to false, you can skip</p>	optional	<code>http://monitoring-prometheus-prometheus.monitoring:9090</code>	string, address		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	this option.					
gateway.scaling.pollingInterval	KEDA polling interval (in seconds) - the interval to check on. See KEDA documentation for more information.	optional	30	integer		
gateway.scaling.maxReplicaCount	Maximum number of replicas that are raised by KEDA/HPA. See KEDA documentation for more information.	optional	100	integer		
gateway.scaling.thresholdSignin	In persons - number of registered agents that causes the Gateway auto-scaling if exceeded	optional	71	integer		
gateway.budget.minAvailable	Option to configure the PodDisruptionBudget option. Do not specify this option if you do not need the PodDisruptionBudget option for the Gateway deployment.	optional		Kubernetes PodDisruptionBudget (PBD) value		gateway: budget:  minAvailable: 50%
gateway.secrets.type	Describes where the secrets are taken - in Kubernetes secrets, CSI driver, or from the Environment	mandatory		csi k8s env		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	variables					
gateway.secrets.csi.gws	If the secrets.type option is set to csi, the name of the CSI object contains the GWS secret			string		
gateway.secrets.k8s.gws	If the secrets.type option is set to k8s, the name of the Kubernetes Secret object that contains the GWS secret			string		
gateway.secrets.env.gwsClient	If the secrets.type option is set to env, the value is GWS clientid created for WebRTC			string		
gateway.secrets.env.gwsSecret	If the secrets.type option is set to env, the value is GWS secret for the client given clientid			string		
gateway.securityContext	Security context for the Gateway container	optional		Specification		gateway: securityContext:  runAsUser: 500  runAsGroup: 500
gateway.serviceAccountName	Name of the ServiceAccount that is used to run the Gateway pod	optional		string		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
gateway.logPath	<p>Path to the log-directory. used for both - PVC or HostPath types of logs. Also, check the esServer option. If /mnt/log/webrtc is specified, the /mnt/log/webrtc//webrtcgw logfiles are created and used in the mentioned path. If the image.initContainerImage option is not specified, the folder with the pod name will not be created and the /mnt/log/webrtc/webrtcgw logfiles will be created.</p> <p><b>Important</b> If this option is set to stdout, the entire WebRTC GW logs are produced to the stdout in JSON format.</p>	mandatory	"/mnt/log/webrtc"	string		"/export/vol1/PAT/infra/webrtc"
gateway.logPvc	Option for Persistent Volume Claim used for the Gateway logs. If logPvc is not defined, the	optional		Section		<pre>gateway:   logPvc:  pvcName: webrtc- gateway- log-pvc</pre>

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	HostPath is used for the logs mount.					<pre> volumeName: webrtc- gateway- log-volume  storageClassName: genesys- webrtc  capacity: 5Gi  volumeSpec:  accessModes: - ReadWriteMany  persistentVolumeReclaimPo Retain nfs:  path: /export/ voll/PAT/ infra/ webrtc  server: 192.168.30.51 </pre>
gateway.logPvc.pvcName	Name of the Persistent Volume Claim. If this option is present, the PVC is created. Else, the hostpath is used for the Gateway logs.	optional		string		
gateway.logPvc.volumeName	PersistentVolume name for the PVC. Single Volume is used for both green and blue deployments of the gateway	optional		string		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
gateway.logPvc.persistentVolumeSpec	If the Persistent Volume specification is configured in the gateway.logPvc.volumeSpec option, the PersistentVolume object with name from the gateway.logPvc.volumeName option is created using this specification.	optional		Specification		<pre>gateway:   logPvc:  volumeSpec:  accessModes: - ReadWriteMany  persistentVolumeReclaimPolicy: Retain  nfs:  path: /export/ voll/PAT/ infra/ webrtc  server: 192.168.30.51</pre>
gateway.logPvc.volumeAnnotations	Any additional annotations that are used for the PersistentVolume if the gateway.logPvc.volumeSpec is specified here.	optional		Specification		<pre>gateway:   logPvc:  volumeAnnotations:  pv.kubernetes.io/ bound-by- controller: 'yes'</pre>
gateway.esServerAddress	Specifies the destination for the ElasticSearch logging - ElasticSearch server address or stdout. Gateway produces messages in the ElasticSearch format.	optional	stdout	network address or "stdout"		
gateway.restartPolicy	Restart policy for gateway pods.	Optional	Always	depends on cluster		
coturn.port	Coturn port that is used by the	optional	443	integer		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	CoTurn Load Balancer					
coturn.lblpBlue	External IP for CoTurn blue Load Balancer service. The IP must be same as the one used for the gateway.turnExternalUriBlue A-record	mandatory		IP address		
coturn.lblpGreen	External IP for CoTurn green Load Balancer service. The IP must be same as the one used for the gateway.turnExternalUriGreen A-record	mandatory		IP address		
coturn.replicas	Number of CoTurn pods	optional	1	integer		
coturn.podAnnotations	Any additional annotations that are applied for CoTurn pods	optional		Specification		<pre> coturn:   podAnnotations:     pods/       realtime:         "true"     pods/       owner:         "1051" </pre>
coturn.resources	Describes resources requested for the CoTurn pods. Do not specify this option if you do not need resources requests/limits.	optional		Section		<pre> coturn:   resources:     requests:       cpu:         "0.5"     memory:       "768Mi"     limits:       memory:         "8Gi" </pre>
coturn.resources.requests.cpu	Requested CPU	optional	0.5	Kubernetes		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	amount of CPU. Coturn requires 0.08CPU per call.			CPU request format		
coturn.resource.requests.memory	Requested amount of memory	optional	150	Kubernetes memory request format		
coturn.resource.limits.memory	Absolute value for the container's memory usage limit	optional	"8Gi"	Kubernetes value for resource limit		
coturn.scaling	Describes the autoscaling parameters. If the deployment . keda option is set to false, you can skip this section	optional		Section		<pre> coturn:   scaling:      pollingInterval:       30      maxReplicaCount:       100      thresholdCpu:       60      thresholdMemory:       60 </pre>
coturn.scaling.pollingInterval	Specifies the KEDA polling interval in seconds - the interval to check each trigger on. Refer to KEDA documentation for more information.	optional	30	integer		
coturn.scaling.maxReplicaCount	Maximum number of replicas that are raised by KEDA/HPA. Refer to KEDA documentation for more information.	optional	100	integer		
coturn.scaling.thresholdSignin		optional	71	integer		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	percentage					
coturn.scaling.thresholdCpu	In percentage. The target value is the average of the CPU resource metric across all pods, represented as a percentage of the requested value of the resource for the pods.	optional	60	integer		
coturn.scaling.thresholdMemory	In percentage. The target value is the average of the memory resource metric across all pods, represented as a percentage of the requested value of the resource for the pods.	optional	60	integer		
coturn.budget.minAvailable	Option to configure PodDisruptionBudget. Do not specify this option, if you do not need PodDisruptionBudget for the CoTurn deployment.	optional		Kubernetes PDB value		coturn: budget:  minAvailable: 50%

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
coturn.securityContext	Security context for the CoTurn container.	optional		Specification		<pre> coturn: securityContext:  runAsUser: 500  runAsGroup: 500 </pre>
coturn.serviceAccountName	Name of the ServiceAccount to use for the CoTurn pod.	optional		string		
coturn.logPath	Path to the log-directory. This can be the directory path or "stdout". This path is used for both PVC or HostPath types of logs. Example: If /mnt/log/webrtc is specified, "/mnt/log/webrtc//turn.xxx.log" logfile is created and used in the mentioned path. If image.initContainerImage is not specified, the folder with pod name will not be created and mnt/log/webrtc/turn.xxx.log logfile will be created.	mandatory	"/mnt/log/webrtc"	string		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
coturn.logPvc	Section for Persistent Volume Claim used for CoTurn logs. If this option not defined, the HostPath is used for logs mount.	optional	"/mnt/log/webrtc"	Section		<pre> coturn:   logPvc:    pvcName:   webrtc-   coturn-log-   pvc    storageClassName:   default    capacity:   10Gi    volumeName:   webrtc-   coturn-log-   volume    volumeSpec:     nfs:    server:   192.168.1.5    path:   /storage/   webrtc    volumeMode:   Filesystem    persistentVolumeReclaimPolicy:   Retain </pre>
coturn.logPvc.pvcName	Name of PersistentVolumeClaim. If this option is present, PVC will be created. Else, the HostPath is used for CoTurn logs.	optional		string		
coturn.logPvc.storageClassName	StorageClass name for CoTurn PVC	optional		string		
coturn.logPvc.capacity	Volume capacity	optional		Kubernetes capacity storage values		
coturn.logPvc.volumeName	Persistent Volume name for the	optional		string		

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	PVC. Single Volume is used for both green and blue deployments of the CoTurn logs					
coturn.logPvc.persistentVolumeSpec	If the Persistent Volume specification is configured in coturn.logPvc.volumeSpec, the Persistent Volume object with name from the coturn.logPvc.volumeName will be created using this specification.	optional		Specification		<pre>gateway:   logPvc:  volumeSpec:  accessModes: -   ReadWriteMany  persistentVolumeReclaimPolicy:   Retain  nfs:  path: /export/ voll/PAT/ infra/ webrtc  server: 192.168.30.51</pre>
coturn.logPvc.persistentVolumeAnnotations	Any additional annotations that are used for the Persistent Volume, if the coturn.logPvc.volumeSpec option is specified	optional		Specification		<pre>gateway:   logPvc:  volumeAnnotations:  pv.kubernetes.io/ bound-by- controller: 'yes'</pre>
coturn.restartPolicy	Restart policy for coturn pods.	optional	Always	depends on cluster		
labels.common	Describes the additional labels for common resources	optional				
labels.gateway	Describes the additional	optional				

Option name	Description	Is mandatory	Default value	Valid value	Notes	Example
	labels for the Gateway resources - pods, deployments, and services					
labels.coturn	Describes the additional labels for the CoTurn resources - pods, deployments, and services	optional				
labels.alerts	Describes the additional labels for the alert objects	optional				

## Configure Kubernetes

*Document the layouts for the following so customers can create them if their Helm chart doesn't include a way to do this:*

- *ConfigMaps*
- *Secrets*

## Configure security

The security context settings define the privilege and access control settings for pods and containers.

By default, the user and group IDs are set in the **values.yaml** file as 500:500:500, meaning the **genesys** user.

```
securityContext:
  runAsNonRoot: true
  runAsUser: 500
  runAsGroup: 500
  fsGroup: 500
```

---

## Arbitrary UIDs in AKS

If you want to use arbitrary UIDs in your Azure Kubernetes Services deployment, override the **securityContext** settings in the **values.yaml** file, so that you do not define any specific IDs.

```
podSecurityContext:
  runAsNonRoot: true
  runAsUser: null
  runAsGroup: 0
  fsGroup: null

securityContext:
  runAsNonRoot: true
  runAsUser: null
  runAsGroup: 0
```

## Configure the service

Before proceeding with the deployment process, perform the following pre-steps:

1. **Review values-template.yaml in helm charts:** It provides all the available options with comments and explanations.
2. **Configure all the options in your own values file:** Configure/overwrite values for options that you need. Use the values-template.yaml file from the package that displays the list of available options with their description.

### Important

Do not configure **deployment.type** and **deployment.color** options in values.yaml-file(s). These values should be used only during deployment process as command-line parameters to specify the deployment process.

Sample values.yaml file:

```
deployment:
  namespace:          webrtc
  ingress:
    domain: apps.vce-c0.eps.genesys.com
    annotations:
      kubernetes.io/ingress.class:      nginx01-internal
      nginx.ingress.kubernetes.io/affinity:      cookie
      nginx.ingress.kubernetes.io/affinity-mode: persistent
      nginx.ingress.kubernetes.io/ssl-redirect: "false"
      nginx.ingress.kubernetes.io/session-cookie-path:  "/; Secure"
      nginx.ingress.kubernetes.io/session-cookie-samesite: None
  dnsPolicy: ClusterFirst
  dnsConfig:
    options:
```

---

```
- name: ndots
  value: "3"
keda: false
coturnDeployment: external

monitoring:
  enabled: false
  dashboards: false
  prometheusMetrics: false
  prometheusAlerts: false

image:
  imagePullSecrets:
  - webrtcjfrogsecret
  initContainerImage: pureengage-docker-staging.jfrog.io/alpine:3.7-curl
  webrtc: pureengage-docker-staging.jfrog.io/webrtc
  coturn: pureengage-docker-staging.jfrog.io/webrtc
  webrtcVersion: 9.0.000.88
  coturnVersion: 9.0.000.88

gateway:
  logPath: "/export/voll/PAT/infra/webrtc"
  logPvc:
    pvcName: webrtc-gateway-log-pvc
    volumeName: webrtc-gateway-log-volume
    storageClassName: genesys-webrtc
    capacity: 5Gi
    volumespec:
      accessModes:
      - ReadWriteMany
      persistentVolumeReclaimPolicy: Retain
      nfs:
        path: /export/voll/PAT/infra/webrtc
        server: 192.168.30.51
  esServer: stdout
  replicas: 1
  workersCount: 1
  voiceSipProxy: voice-siproxy.voice.svc.cluster.local:5080;transport=tcp
  turnExternalUriBlue: 192.168.30.208
  turnExternalUriGreen: 192.168.30.209
  authRedirectUri: http://gauth.apps.vce-c0.eps.genesys.com:80
  authService: http://gauth-auth.gauth.svc.cluster.local:80
  envService: https://gws.apps.vce-c0.eps.genesys.com
  resources:
    requests:
      # NB! 800m per worker, MUST be integer, not string - will be multiplied by
workersCount in helm
      cpu: 800
      # NB! 150Mi per worker, MUST be integer, not string - will be multiplied by
workersCount in helm
      memory: 150
    limits:
      memory: "8Gi"
  secrets:
    type: env
    env:
      gwsClient: external_api_client
      gwsSecret: secret
  securityContext:
    runAsUser: 500
```

---

---

```
runAsGroup: 500

coturn:
  logPath: "/export/vol1/PAT/infra/coturn/"
  logPvc:
    pvcName: webrtc-coturn-log-pvc
    volumeName: webrtc-coturn-log-volume
    storageClassName: genesys-webrtc
    capacity: 5Gi
    volumeSpec:
      accessModes:
        - ReadWriteMany
      persistentVolumeReclaimPolicy: Retain
      nfs:
        path: /export/vol1/PAT/infra/webrtc
        server: 192.168.30.51
  replicas: 1
  port: 443
  lbIpBlue: 192.168.30.208
  lbIpGreen: 192.168.30.209
  securityContext:
    runAsUser: 500
    runAsGroup: 500
```

3. **PersistentVolume (PV) and PersistentVolumeClaim (PVC):** If you plan to use PV for logs, create the PV and then specify it for PVC of Gateway and CoTurn.

PV can also be created during the common-infrastructure deployment. You should review the values-template.yaml file and then configure the PV specification for Gateway and CoTurn.

Single PV/PVC pair will be used for both Green and Blue deployments of Gateway, and another single PV/PVC pair will be used for both Green and Blue deployments of CoTurn.