

# CSCF Scaling Management

## Call Session Control Function

---

### USER GUIDE

**Copyright**

© Ericsson AB 2017, 2018. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

**Disclaimer**

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

**Trademark List**

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Prerequisites	1
<b>2</b>	<b>Description</b>	<b>3</b>
2.1	Auto Scale-Out	3
2.2	Graceful Scale-In	3
2.3	Forceful Scale-In	4
<b>3</b>	<b>Prepare for Scaling</b>	<b>7</b>
<b>4</b>	<b>Manage Scaling Manually</b>	<b>9</b>
4.1	Configure Scale-Out Manually	9
4.2	Configure Graceful Scale-In Manually	10
4.3	Configure Forceful Scale-In Manually	13
<b>5</b>	<b>Manage Scaling from Cloud with Heat Orchestration</b>	<b>17</b>
5.1	Configure Scale-Out with Heat Orchestration	17
5.2	Configure Graceful Scale-In with Heat Orchestration	18
5.3	Configure Forceful Scale-In with Heat Orchestration	24
<b>6</b>	<b>Manage Scaling with VNF-LCM</b>	<b>27</b>





# 1 Introduction

This document describes the scalability functions of the Call Session Control Function (CSCF) cluster as a distributed system. It also gives instructions on how to do expansion or contraction of the cluster using these functions.

If the scaling type is not mentioned, this document always refers to horizontal scaling, where the scalability of the system is provided by multiple instances to distribute the load in parallel for having the capacity needed. Vertical scaling is not considered in this document.

## 1.1 Prerequisites

This section describes the prerequisites that must be fulfilled before expanding or contracting the CSCF cluster.

### 1.1.1 Licenses

The scaling function does not require a license.

### 1.1.2 Documents

Before starting these procedures, the following documents must be available:

- *CSCF Health Check*
- *Ericsson Command-Line Interface User Guide*

### 1.1.3 Prerequisites

Before starting this procedure, ensure that the following conditions are met:

- The procedure must only be performed by support personnel with experience of Cloud and the CSCF.
- No other upgrade or maintenance activity must be performed during the procedure.
- Before starting these procedures, the user performing the operations must have access to the System Controller (SC) nodes.
- Signaling Manager Command-Line Interface (CLI) or Graphical User Interface (GUI) must be closed before the start of the Scaling Operations. Manual updates of the configurations during Scaling Operations are not allowed.



- A Virtual Infrastructure Manager (VIM) is available.
- For scaling operations without the VNF Lifecycle Manager (VNF-LCM), an Ericsson Command-Line Interface (ECLI) session in `EXEC` mode is in progress.


#### **1.1.4 Limitations**

Even though the PL-3 and PL-4 Virtual Machines (VMs) are considered to be part of the scaling domain, they cannot be scaled in.


## 2 Description

### 2.1 Auto Scale-Out

Auto Scale-Out is an operation where one or more new compute resources are launched, see Figure 1. The system automatically detects, configures, and brings up the nodes as a member of the scaling domain of the cluster. See Figure 2 for an example where one new compute node is added to the cluster.



*Figure 1 A New Compute Resource Is Spawned and Available*



*Figure 2 After Auto Scale-Out, a New Resource Is Added to the Cluster*

### 2.2 Graceful Scale-In

Graceful Scale-In is an operation where one or more compute resources, part of the scaling domain of the cluster (see Figure 3) are removed from the cluster (see Figure 4) to free up resources.




Figure 3 The Node Named PL-(N-1) Is Part of the Cluster




Figure 4 The Node Named PL-(N-1) Is Removed from the Cluster and Its Resources Can Be Released

**Note:** The Graceful Scale-In operation can be rejected by the cluster if, according to the automatic estimation of the system, the target size of the cluster does not have the memory resources to serve the needed memory capabilities for the ongoing traffic.

## 2.3 Forceful Scale-In

Forceful Scale-In is, similarly to Graceful Scale-In, an operation to remove one or more nodes from the scaling domain of the cluster. The only difference is that in this case, either the node is not available (see Figure 5) or scale-in with potential traffic loss is acceptable. If the node is not available, it can be either because it already freed up its resources or because of a failure. Therefore the removal is only an administrative operation, see Figure 6.





Figure 5 The Node Named PL-(N-1) in the Cluster Scaling Domain Is Unavailable





**Figure 6** *The Node Named PL-(N-1) Is Removed Administratively from the Cluster*





### 3 Prepare for Scaling

Before any scaling starts, make sure that the scaling feature is enabled and a system backup is created.

1. Connect to one of the SC nodes:

```
ssh <user>@<system management IP address>
```

2. Check the operational state of the scaling feature:

```
SC-1: ~ # cmw-configuration --status SCALING
```

The following is an example output:

```
Disable
```

3. If the result is `Enable`, scaling is prepared. Exit this procedure.

4. If the result is `Disable`, enable scaling functionality:

```
SC-1: ~ # cmw-configuration --enable SCALING
```

5. Before any scaling-related activities are performed, create a system backup. Refer to *Create Backup*.





## 4 Manage Scaling Manually

Before any scaling starts, make sure that the scaling feature is enabled and a system backup is created. See Section 3 on page 7.

**Note:** Do not use this procedure, if the CSCF was instantiated using Heat or VNF-LCM.

### 4.1 Configure Scale-Out Manually

1. Check that the cluster is in a healthy state, refer to *CSCF Health Check*.
2. Create VMs that must have the same number of Virtual CPUs (vCPUs), the same amount of RAM, and the same number of ports as the other Payload (PL) VMs in the cluster. Refer to the VIM documentation for how to create VMs.
3. Launch the new VMs in the VIM. The new VMs automatically PXE boot from the System Controller (SC) node VMs.
4. Navigate to the *CrM* Managed Object (MO), for example:

```
>dn ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1
```

5. Verify that the scale-out process has started:

```
(CrM=1) >show -r
```

The following is an example output:

```
CrM=1
autoRoleAssignment=ENABLED
ComputeResourceRole=PL-3
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,Role=Default-Role"
  uses="ManagedElement=1,Equipment=1,ComputeResource=PL-3"
ComputeResourceRole=PL-4
  adminState=UNLOCKED
  instantiationState=INSTANTIATING
  operationalState=DISABLED
  provides="ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,Role=Default-Role"
  uses="ManagedElement=1,Equipment=1,ComputeResource=PL-4"
Role=SYSTEM
  isProvidedBy
  scalability=NON_SCALABLE
Role=Default-Role
  isProvidedBy
    "ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,ComputeResourceRole=PL-3"
    "ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,ComputeResourceRole=PL-4"
  scalability=SCALABLE
```

6. Continue to check the progress until the scale-out process has ended and that the added node has joined the cluster:

```
(CrM=1) >show -m ComputeResourceRole -p \  
instantiationState,operationalState
```

The following example output shows the final result:

```
ComputeResourceRole=PL-3  
  instantiationState=INSTANTIATED  
  operationalState=ENABLED  
ComputeResourceRole=PL-4  
  instantiationState=INSTANTIATED  
  operationalState=ENABLED
```

This example shows that `instantiationState` has changed to `INSTANTIATED` for node PL-4. It means that PL-4 is added to the cluster.

The example also shows that `operationalState` has changed to `ENABLED` for node PL-4. It means that node PL-4 has joined the cluster.

7. Perform a health check, refer to *CSCF Health Check*.

## 4.2 Configure Graceful Scale-In Manually

1. Check that the cluster is in a healthy state, refer to *CSCF Health Check*.
2. Navigate to the *CrM* MO, for example:

```
>dn ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1
```

3. Verify that the VM to be scaled-in is scalable, for example:

```
(CrM=1) >show -r
```



```
CrM=1
autoRoleAssignment=ENABLED
ComputeResourceRole=PL-5
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
ComputeResourceRole=PL-7
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-7"
ComputeResourceRole=PL-4
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-4"
ComputeResourceRole=PL-3
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-3"
ComputeResourceRole=SC-2
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
  uses="ManagedElement=1, Equipment=1, ComputeResource=SC-2"
ComputeResourceRole=SC-1
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
  uses="ManagedElement=1, Equipment=1, ComputeResource=SC-1"
ComputeResourceRole=PL-6
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
Role=SYSTEM
  isProvidedBy
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=SC-1"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=SC-2"
  scalability=NON_SCALABLE
Role=Default-Role
  isProvidedBy
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=PL-6"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=PL-3"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=PL-4"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=PL-5"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResource=PL-7"
  scalability=SCALABLE
```

4. Retrieve the Universally Unique Identifier (UUID) for the VM to be scaled-in, for example:

```
(CrM=1) >show ManagedElement=1, Equipment=1, \
ComputeResource=PL-8
```

The following is an example output:



```
ComputeResource=PL-8
  macAddress
    "fa:16:3e:b7:d3:a3"
    "fa:16:3e:27:cb:90"
    "fa:16:3e:24:73:4a"
  uuid="a4dcda89-cf95-4bf6-81bd-99d47fde9eef"
```

5. Navigate to the *ComputeResourceRole* MO for the VM to be scaled-in, for example:

```
(CrM=1) >ComputeResourceRole=PL-8
```

6. Enter Config mode:

```
(ComputeResourceRole=PL-8) >configure
```

7. Prepare the scale-in operation, for example:

```
(config-ComputeResourceRole=PL-8) >no provides
```

8. Navigate to the *CrM* MO:

```
(config-ComputeResourceRole=PL-8) >up
```

9. Perform the scale-in:

```
(config-CrM=1) >commit
```

**Note:** To cancel the scale-in, run **abort**.

10. Verify that the scale-in process has started, for example:

```
(CrM=1) >show -r
```

11. If a failure occurs during the scale-in, refer to *CSCF Troubleshooting Guideline*.

12. Verify that the VM is scaled-in:

```
(CrM=1) >show -r
```

The following is an example output showing that the VM *ComputeResourceRole=PL-8* is no longer running:





```
CrM=1
autoRoleAssignment=ENABLED
ComputeResourceRole=PL-3
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-3"
ComputeResourceRole=PL-4
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-4"
ComputeResourceRole=SC-1
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
  uses="ManagedElement=1, Equipment=1, ComputeResource=SC-1"
ComputeResourceRole=SC-2
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
  uses="ManagedElement=1, Equipment=1, ComputeResource=SC-2"
ComputeResourceRole=PL-5
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
ComputeResourceRole=PL-6
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
ComputeResourceRole=PL-7
  adminState=UNLOCKED
  instantiationState=INSTANTIATED
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-7"
Role=Default-Role
  isProvidedBy
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-4"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-5"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-6"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-7"
  scalability=SCALABLE
Role=SYSTEM
  isProvidedBy
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-3"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-1"
    "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-2"
  scalability=NON_SCALABLE
```

13. Remove the VM with the UUID that is retrieved in Step 4 from the VIM. Refer to the VIM documentation.

14. Perform a health check, refer to *CSCF Health Check*.

## 4.3 Configure Forceful Scale-In Manually

1. Check that the cluster is in a healthy state, refer to *CSCF Health Check*.
2. Remove one of the scalable VMs from the VIM.

**Note:** Do NOT delete any of the VMs named SC-1, SC-2, PL-3, or PL-4.



3. Navigate to the *CrM* MO, for example:

```
>dn ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1
```

4. Identify the *ComputeResourceRole* where the *adminState* is *LOCKED* and *operationalState* is *DISABLED*, for example:

```
(CrM=1) >show -r
```

```
ComputeResourceRole=PL-8
adminState=LOCKED
instantiationState=INSTANTIATED
operationalState=DISABLED
provides="ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,Role=Default-Role"
uses="ManagedElement=1,Equipment=1,ComputeResource=PL-8"
```

5. Enter Config mode:

```
(ComputeResourceRole=PL-8) >configure
```

6. Prepare the scale-in operation, for example:

```
(config-ComputeResourceRole=PL-8) >no provides
```

7. Navigate to the *CrM* MO:

```
(config-ComputeResourceRole=PL-8) >up
```

8. Perform the scale-in:

```
(config-CrM=1) >commit
```

9. Verify that the scaling-in process has started, for example:

```
(CrM=1) >show -r
```

The following is an example output:

```
ComputeResourceRole=PL-8
adminState=LOCKED
instantiationState=INSTANTIATED
operationalState=DISABLED
provides="ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,Role=Default-Role"
uses="ManagedElement=1,Equipment=1,ComputeResource=PL-8"
```

10. If a failure occurs during the scale-in, refer to *CSCF Troubleshooting Guideline*.

11. Verify that the VM is scaled-in:

```
(CrM=1) >show -r
```

The following is an example output showing that the VM *ComputeResourceRole=PL-8* is no longer running:



```

CrM=1
  autoRoleAssignment=ENABLED
  ComputeResourceRole=PL-3
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-3"
  ComputeResourceRole=PL-4
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-4"
  ComputeResourceRole=SC-1
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-1"
  ComputeResourceRole=SC-2
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-2"
  ComputeResourceRole=PL-5
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
  ComputeResourceRole=PL-6
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
  ComputeResourceRole=PL-7
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-7"
  Role=Default-Role
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-4"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-5"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-6"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-7"
    scalability=SCALABLE
  Role=SYSTEM
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-3"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-1"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-2"
    scalability=NON_SCALABLE

```

## 12. Perform a health check, refer to *CSCF Health Check*.





## 5 Manage Scaling from Cloud with Heat Orchestration

Before scaling the VNF on the cloud through Heat orchestration, make sure that the scaling feature is enabled and the VNF is instantiated with Heat Orchestration Template (HOT) that support scaling.

**Note:** Do not use this procedure, if the VNF was instantiated using VNF-LCM or manually.

### 5.1 Configure Scale-Out with Heat Orchestration

1. Check that the cluster is in a healthy state, refer to *CSCF Health Check*.

2. Check that the status of the CSCF stack is `CREATE_COMPLETE` or `UPDATE_COMPLETE`:

```
heat stack-list
```

If the status of the stack is not `CREATE_COMPLETE` or `UPDATE_COMPLETE`, stop the scaling procedure. For information on how to identify and correct the stack status, refer to VIM documentation.

3. Check the value of parameter `number_of_scaled_out_PL_VMs`.

```
heat stack-show <CSCF stack name> | \
grep number_of_scaled_out_PL_VMs
```

4. Increase the value of parameter `number_of_scaled_out_PL_VMs` by the number of VMs to be scaled out.

For example: The current value of the parameter `number_of_scaled_out_PL_VMs` is 1 (meaning: beyond the initial size of 2+2; the cluster contains an extra VM/PL, so the size of the VNF is actually 2+3). To increase the size of the cluster to 2+5, that is, scale out by 2 VMs, the new value of the parameter should be 3.

5. Update the stack:

```
heat stack-update <CSCF stack name> -x -P \
number_of_scaled_out_PL_VMs=<number_of_scaled_out_PL_VMs>
```

6. Monitor the progress of the stack-update until the stack status is `UPDATE_COMPLETE`:

```
heat stack-list
```

7. If the stack status is not `UPDATE_COMPLETE`, check the reason and Troubleshoot the issue as described in *CSCF Troubleshooting Guideline* and then repeat Step 5:

```
heat stack-show
```

8. Navigate to the *CrM* MO, for example:

```
>dn ManagedElement=NODE06ST,SystemFunctions=1,\  
SysM=1,CrM=1
```

9. Verify that the new VMs are added and enabled, for example:

```
(CrM=1) >show -r
```

**Note:** It takes a few minutes until the VMs added in the stack shows up in the *CrM* MO.

```
CrM=1  
autoRoleAssignment=ENABLED  
ComputeResourceRole=PL-5  
adminState=UNLOCKED  
instantiationState=INSTANTIATED  
operationalState=ENABLED  
provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"  
uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
```

10. Perform a health check, refer to *CSCF Health Check*.

## 5.2 Configure Graceful Scale-In with Heat Orchestration



---

### Attention!

Risk of data loss or data corruption.

---

Do not use this procedure to remove resources created manually by Heat commands (**nova**, **neutron**), or from Horizon or Atlas Dashboard as it can corrupt the database of Heat. To repair a faulty resource of a heat stack, use the following Heat commands on the stack: `heat action-check`, `heat resource-list`, and `heat stack-update`. Detailed descriptions of these procedures are beyond the scope of this instruction. For more information about Heat, refer to <https://wiki.openstack.org/wiki/Heat>.

1. Check that the status of the CSCF stack is `CREATE_COMPLETE` or `UPDATE_COMPLETE`:

```
heat stack-list
```



If the status of the stack is not `CREATE_COMPLETE` or `UPDATE_COMPLETE`, stop the scaling procedure. For information on how to identify and correct the stack status, refer to VIM documentation.

2. Check that the cluster is in a healthy state, refer to *CSCF Health Check*.
3. Navigate to the *CrM* MO, for example:

```
>dn ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1
```

4. Verify that the VM to be scaled-in is scalable, for example:

```
(CrM=1) >show -r
```



```

CrM=1
  autoRoleAssignment=ENABLED
  ComputeResourceRole=PL-8
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-8"
  ComputeResourceRole=PL-5
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
  ComputeResourceRole=PL-7
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-7"
  ComputeResourceRole=PL-4
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-4"
  ComputeResourceRole=PL-3
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-3"
  ComputeResourceRole=SC-2
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-2"
  ComputeResourceRole=SC-1
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-1"
  ComputeResourceRole=PL-6
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
  Role=SYSTEM
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-1"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-2"
    scalability=NON_SCALABLE
  Role=Default-Role
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-6"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-3"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-4"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-5"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-7"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-8"
    scalability=SCALABLE

```

5. Retrieve the Universally Unique Identifier (UUID) for the VM to be scaled-in and write it down for use in Step a, for example:

```

(CrM=1) > show ManagedElement=1, Equipment=1, \
ComputeResource=PL-8

```





```
ComputeResource=PL-8
  macAddress
    "fa:16:3e:b7:d3:a3"
    "fa:16:3e:27:cb:90"
    "fa:16:3e:24:73:4a"
  uuid="a4dcda89-cf95-4bf6-81bd-99d47fde9eef"
```

6. Navigate to the *ComputeResourceRole* MO for the VM to be scaled-in, for example:

```
(CrM=1) >ComputeResourceRole=PL-8
```

7. Enter Config mode:

```
(ComputeResourceRole=PL-8) >configure
```

8. Prepare the scale-in operation, for example:

```
(config-ComputeResourceRole=PL-8) >no provides
```

9. Navigate to the *CrM* MO:

```
(config-ComputeResourceRole=PL-8) >up
```

10. Perform the scale-in:

```
(config-CrM=1) >commit
```

**Note:** To cancel the scale-in, run **abort**.

11. Verify that the scaling-in process has started, for example:

```
(CrM=1) >show -r
```

12. If a failure occurs during the scale-in, refer to *CSCF Troubleshooting Guideline*.

13. Verify that the VM is scaled-in:

```
(CrM=1) >show -r
```

The following is an example output showing that the VM *ComputeResourceRole=PL-8* is no longer running:



```

CrM=1
  autoRoleAssignment=ENABLED
  ComputeResourceRole=PL-3
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-3"
  ComputeResourceRole=PL-4
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-4"
  ComputeResourceRole=SC-1
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-1"
  ComputeResourceRole=SC-2
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-2"
  ComputeResourceRole=PL-5
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
  ComputeResourceRole=PL-6
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
  ComputeResourceRole=PL-7
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-7"
  Role=Default-Role
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-4"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-5"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-6"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-7"
    scalability=SCALABLE
  Role=SYSTEM
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-3"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-1"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-2"
    scalability=NON_SCALABLE

```

14. If more VMs need scale-in, repeat Step 5 to Step 13

15. Check the value of parameter `number_of_scaled_out_PL_VMs`.

```

heat stack-show <CSCF stack name> | \
grep number_of_scaled_out_PL_VMs

```

```

"number_of_scaled_out_PL_VMs": "3"

```

16. Find the value for parameter `PL_to_be_scaled_in`:



- a. Check the UUID that is retrieved in Step 5, and write down its position.

```
heat stack-show <CSCF stack name> | grep <uuid>
```

The following example shows that the position of the example UUID `a4dcda89-cf95-4bf6-81bd-99d47fde9eef` that is retrieved in Step 5 is the second in the list:

```
"output_value": "55ab4245-8b18-49ca-8ded-1cfca49a4d89 \
a4dcda89-cf95-4bf6-81bd-99d47fde9eef baa95296-b433-4801-a2fc-d6ea23520287 \
9ab5e3a4-aa31-4bc4-8e0a-d31f7f28c138"
```

- b. Retrieve the names of the scaled-out VMs from the stack:

```
heat stack-show <CSCF stack name> | \
grep <CSCF stack name>_scaled_out_VM
```

```
"output_value": "cscf-104_scaled_out_VM-0 cscf-104_scaled_out_VM-9 \
cscf-104_scaled_out_VM-10 cscf-104_scaled_out_VM-11"
```

- c. Find the name of the VM to be scaled in at the position from Step a in the list from Step b, and write down the number at the end of the VM name.

The number is the index of the VM and is used as the value of parameter `PL_to_be_scaled_in`. For example, write down 9 that is the number at the end of the VM names in the example from Step b.

**Note:** The index starts with 0.

17. If multiple VMs are scaled in, repeat Step 16.

18. Update the stack:

```
heat stack-update <CSCF stack name> -x -P \
number_of_scaled_out_PL_VMs=<number_of_scaled_out_PL_VMs> \
-P PL_to_be_scaled_in=<index of VMs>
```

The value for the parameter `number_of_scaled_out_PL_VMs` is the value obtained in Step 15 minus the number of VMs to scale in. The value for the parameter `PL_to_be_scaled_in` is a comma-separated list of the values obtained in Step 16 (repeat Step 16 for multiple VMs to scale in).

According to the example in Step a and Step b, the parameter `PL_to_be_scaled_in` is to set to 9.

## 5.3 Configure Forceful Scale-In with Heat Orchestration



### Attention!

Risk of data loss or data corruption.

Do not use this procedure to remove resources created manually by Heat commands (**nova**, **neutron**), or from Horizon or Atlas Dashboard as it can corrupt the database of Heat. To repair a faulty resource of a heat stack, use the following Heat commands on the stack: **heat action-check**, **heat resource-list**, and **heat stack-update**. Detailed descriptions of these procedures are beyond the scope of this instruction. For more information about Heat, refer to <https://wiki.openstack.org/wiki/Heat>.

1. Check that the status of the CSCF stack is **CREATE\_COMPLETE** or **UPDATE\_COMPLETE**:

```
heat stack-list
```

If the status of the stack is not **CREATE\_COMPLETE** or **UPDATE\_COMPLETE**, stop the scaling procedure. For information on how to identify and correct the stack status, refer to VIM documentation.

2. Check that the cluster is in a healthy state, refer to *CSCF Health Check*.
3. Check the value of parameter **number\_of\_scaled\_out\_PL\_VMs**.

```
heat stack-show <CSCF stack name> | \
grep number_of_scaled_out_PL_VMs

"number_of_scaled_out_PL_VMs": "3"
```

4. Update the stack:

The value for the parameter **number\_of\_scaled\_out\_PL\_VMs** is the value obtained in Step 3 minus the number of VMs to scale in.

```
heat stack-update <CSCF stack name> -x -P \
number_of_scaled_out_PL_VMs=<number_of_scaled_out_PL_VMs>
```

5. Navigate to the *CrM* MO, for example:

```
>dn ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1
```

6. Identify the **ComputeResourceRole** where the **adminState** is **LOCKED** and **operationalState** is **DISABLED**, for example:

```
(CrM=1)>show -r
```



```
ComputeResourceRole=PL-8
adminState=LOCKED
instantiationState=INSTANTIATED
operationalState=DISABLED
provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
uses="ManagedElement=1, Equipment=1, ComputeResource=PL-8"
```

7. Enter Config mode:

```
(ComputeResourceRole=PL-8) >configure
```

8. Prepare the scale-in operation, for example:

```
(config-ComputeResourceRole=PL-8) >no provides
```

9. Navigate to the *CrM* MO:

```
(config-ComputeResourceRole=PL-8) >up
```

10. Perform the scale-in:

```
(config-CrM=1) >commit
```

11. Verify that the scaling-in process has started, for example:

```
(CrM=1) >show -r
```

The following is an example output:

```
CrM=1
[ ... ]
    ComputeResourceRole=PL-8
    adminState=SHUTTINGDOWN
    instantiationState=UNINSTANTIATING
[ ... ]
```

12. If a failure occurs during the scale-in, refer to *CSCF Troubleshooting Guideline*.

13. Verify that the VM is scaled-in:

```
(CrM=1) >show -r
```

The following is an example output showing that the VM *ComputeResourceRole=PL-8* is no longer running:



```

CrM=1
  autoRoleAssignment=ENABLED
  ComputeResourceRole=PL-3
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-3"
  ComputeResourceRole=PL-4
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-4"
  ComputeResourceRole=SC-1
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-1"
  ComputeResourceRole=SC-2
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=SYSTEM"
    uses="ManagedElement=1, Equipment=1, ComputeResource=SC-2"
  ComputeResourceRole=PL-5
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-5"
  ComputeResourceRole=PL-6
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
  ComputeResourceRole=PL-7
    adminState=UNLOCKED
    instantiationState=INSTANTIATED
    operationalState=ENABLED
    provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
    uses="ManagedElement=1, Equipment=1, ComputeResource=PL-7"
  Role=Default-Role
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-4"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-5"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-6"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-7"
    scalability=SCALABLE
  Role=SYSTEM
    isProvidedBy
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-3"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-1"
      "ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=SC-2"
    scalability=NON_SCALABLE

```

14. If more VMs need scale-in, repeat Step 6 to Step 13.

15. Perform a health check, refer to *CSCF Health Check*.



## 6 Manage Scaling with VNF-LCM

Refer to *CSCF VNF Lifecycle Management*.