

Scale In Cluster Using Heat Orchestration

MTAS

OPERATING INSTRUCTIONS

Copyright

© Ericsson AB 2018, 2019. 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

1	Description	1
2	Procedure	3
2.1	Scale In Cluster Using Heat Orchestration	3
2.2	Prepare for Scaling	4
2.3	Configure Graceful Scale-In with Heat Orchestration	4
2.4	Configure Forceful Scale-In with Heat Orchestration	10





1 Description

This instruction describes how to decrease the MTAS cluster capacity by performing either a graceful or forceful scale-in operation with heat orchestration, which removes a Virtual Machine (VM) from the cluster.

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.

The scaling function does not require a license.

Note: 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 Procedure

2.1 Scale In Cluster Using Heat Orchestration

Prerequisites

- This instruction references the following documents:
 - MTAS Hardening Guide.
 - MTAS Health Check
 - Ericsson Command-Line Interface User Guide
- No tools are required.
- The following conditions must apply:
 - The procedure must only be performed by support personnel with experience of Cloud and MTAS.
 - No other upgrade or maintenance activity must be performed during the procedure.
 - Scaling must only be performed after site-specific initial configuration is applied on the node. For more details, see [MTAS Hardening Guide](#).
 - 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.
 - An Ericsson Command-Line Interface (ECLI) session in Exec mode is in progress.

Steps

1. Prepare for scaling, see Section 2.2 Prepare for Scaling on page 4.
2. Select action based on the type of scale-in:
 - Decrease capacity through graceful scale-in, see Section 2.3 Configure Graceful Scale-In with Heat Orchestration on page 4.
 - Decrease capacity through forceful scale-in, see Section 2.4 Configure Forceful Scale-In with Heat Orchestration on page 10.



2.2 Prepare for Scaling

Steps

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.
See [Create Backup](#).

2.3 Configure Graceful Scale-In with Heat Orchestration



Attention!

Risk of data loss or data corruption.

Do not remove resources created by Heat manually by 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
- heat stack-update

Detailed descriptions of these procedures are beyond the scope of this instruction. For more information about Heat, see <https://wiki.openstack.org/wiki/Heat>.

Steps



1. Make sure that the scaling feature is enabled and a system backup is created, see Section 2.2 Prepare for Scaling on page 4.
2. Check that the status of the MTAS 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 the VIM documentation.

3. Check that the cluster is in a healthy state, see MTAS Health Check.
4. Navigate to the CrM MO via ECLI, for example:

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

5. 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

```

6. 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"
```

7. Navigate to the `ComputeResourceRole` MO for the VM to be scaled in, for example:

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

8. Enter Config mode:

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

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

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

10. Navigate to the `CrM` MO:

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

11. Perform the scale-in:

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

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

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

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

13. If a failure occurs during the scale-in, see [MTAS Troubleshooting Guideline](#).

14. 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

```

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

16. Check the value of parameter `number_of_scaled_out_VMs`.

```

heat stack-show <MTAS stack name> | \
grep number_of_scaled_out_VMs

"number_of_scaled_out_VMs": "3"

```

17. Find the value for parameter `VM_to_be_scaled_in`:

a. Check the UUID that is retrieved in Step 6, and write down its position.



```
heat stack-show <MTAS 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 6 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 <MTAS stack name> | \
grep <MTAS stack name>_scaled_VM
```

The following is an example output:

```
"output_value": "mtas-104_scaled_VM-0 \
mtas-104_scaled_VM-9 mtas-104_scaled_VM-10 \
mtas-104_scaled_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 VM_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.

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

19. Update the stack:

```
heat stack-update <MTAS stack name> -x -P \
number_of_scaled_out_VMs=<number_of_scaled_out_VMs> \
-P VM_to_be_scaled_in=<index of VMs>
```

The value for the parameter number_of_scaled_out_VMs is the value obtained in Step 16 minus the number of VMs to scale in. The value for the parameter VM_to_be_scaled_in is a comma-separated list of the values obtained in Step 17 (repeat Step 17 for multiple VMs to scale in).

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

Note: On certain CEE or Openstack versions the -x parameter might be missing. In this case the heat template file and the environment file must both be provided on the command line. See the documentation of heat for further information.



2.4 Configure Forceful Scale-In with Heat Orchestration



Attention!

Risk of data loss or data corruption.

Do not remove resources created by Heat manually by 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`
- `heat stack-update`

Detailed descriptions of these procedures are beyond the scope of this instruction. For more information about Heat, see <https://wiki.openstack.org/wiki/Heat>.

Steps

1. Check that the status of the MTAS 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 the VIM documentation.

2. Check that the cluster is in a healthy state, see [MTAS Health Check](#).
3. Check the value of parameter `number_of_scaled_out_VMs`.

```
heat stack-show <MTAS stack name> | \
grep number_of_scaled_out_VMs
```

```
"number_of_scaled_out_VMs": "3"
```

4. See Step 17 on how to get the index of VMs to scale in.
5. Update the stack:

The value for the parameter `number_of_scaled_out_VMs` is the value obtained in Step 3 minus the number of VMs to scale in.



```
heat stack-update <MTAS stack name> -x -P \
number_of_scaled_out_VMs=<number_of_scaled_out_VMs> \
-P VM_to_be_scaled_in=<index of VMs>
```

Note: On certain CEE or Openstack versions the `-x` parameter might be missing. In this case the heat template file and the environment file must both be provided on the command line. See the documentation of heat for further information.

6. Navigate to the CrM MO via ECLI, for example:

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

7. 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"
```

8. Enter Config mode:

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

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

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

10. Navigate to the CrM MO:

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

11. Perform the scale-in:

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

12. 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
  [...]
```

13. If a failure occurs during the scale-in, see MTAS Troubleshooting Guideline.



14. 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
```

15. If more VMs need scale-in, repeat Step 7 to Step 14.

16. Perform a health check, see MTAS Health Check.