

IPWorks Scaling Guide for KVM

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 because of 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	Introduction	1
1.1	Prerequisites	1
1.2	Preparation	2
1.3	Related Information	15
2	Scaling Procedure	17
2.1	Creating Backup before Scaling	17
2.2	Scale-Out Operation	17
2.3	Scale-In Operation	20
2.4	Scaling Health Check	25
2.5	Creating the Final Backup	27
	Reference List	29





1 Introduction

This document describes how to perform the scaling operations for IPWorks. Scaling operations can only be performed on 2+2 Standard configuration. IPWorks DHCP or DNS service does not support scalability.

These operations are done in two ways:

- Scale out, with the creation of new Payloads (PLs) by instantiation of new Virtual Machines (VMs), joins them to the cluster and increases the number of PLs in the system.

After scale-out operation, the maximum number of PLs is 10.

- Scale in, with the removal of existing PLs from the cluster and deletion the associated VM instances, decreases the number of PLs in the system.

After scale-in operation, the minimum number of PLs is 2.

1.1 Prerequisites

This section states the prerequisites for performing the scaling procedure. It is assumed that users of this document are familiar with performing operations in their KVM environment.

1.1.1 Prerequisites for IPWorks VNF

The following conditions must apply in IPWorks VNF:

- The scaling functionality is supported starting from IPWorks 1.9.
- The scaling functionality is enabled.

```
cmw-configuration --status SCALING
```

Enable

- No other maintenance activity is in progress in the IPWorks VNF.
- Scheduled backups **MUST** not be active. To disable scheduled backups, refer to **BRF-C Management Guide**.
- IPWorks VNF is in a healthy state, follow Section 2.4 Scaling Health Check on page 25 to check the health status.
- Scaling (both scale-out and scale-in) operation activities are recommended to be performed at off-peak hours.



- eVIP configuration must be based on the IPWorks eVIP template. Any other eVIP configuration is not supported.
- Enough available IP addresses of Signaling network (<IPW_SIG_SP1_NW>) and Data network (<IPW_DATA_SP1_NW>) are reserved for the use of scaled-out PL VM. If there is not enough IP address for scaling-out, expand the Signaling network and Data network.

1.1.2 Conditions

By default, most actions are performed on a host machine, and some actions are performed on Service Controllers (SCs), unless otherwise specified.

1.2 Preparation

Before you perform the scaling operation, the following preparations must be done:

- Check SS7 configuration, refer to Section 1.2.1 Configuring SS7 to Support Scaling Operations on page 2.
- Check eVIP configuration, refer to Section 1.2.2 Configuring eVIP to Support Scaling Operations on page 7.
- Check the tools, refer to Section 1.2.3 Checking the Tools on page 10.
- Check the configuration data, refer to Section 1.2.4 Adapting Configuration File on page 11.

1.2.1 Configuring SS7 to Support Scaling Operations

This section describes how to precheck the SS7 configuration for supporting the scaling function.

Prerequisites:

This section is only applicable when the following conditions are satisfied, if none of the condition is met, skip the section.

- IPWorks uses SS7 stack. The SS7 stack is used by the following IPWorks functions/scenarios:
 - IPWorks ENUM Number Portability (NP)
 - Wi-Fi Mobility Management (WiFi MM)
 - Diameter over SCTP
 - Secure Entitlement Server (SES) Support



– Wi-Fi AAA

Before do the precheck, start the Signaling Manager on the SC-1:

1. Log on to the SC-1.

```
# ssh root@<SC-1 IP address>
```

2. Find the path to PSO storage where SS7 configuration files are stored:

```
# cat /usr/share/pso/storage-paths/config
```

```
<path to config PSO storage>
```

3. Create links to the path where SS7 configuration files are stored, If the path /opt/sign/etc already exists, skip it.

```
# ln -s <path to config PSO storage>/ss7caf-ana90137/etc  
/opt/sign/etc
```

4. Start Signaling Manager on the SC-1.

```
# /opt/sign/EABss7050/bin/signmngui -own.conf  
/opt/sign/etc/signmgr.cnf &
```

Note:

- If the JAVA cannot be found, use the command **export JAVA_HOME=/opt/sign/EABss7069/jre**
- If no X11 DISPLAY variable was set, try to log out the SC-1, and then log on again by using the **-X** option:

```
# ssh -X root@<SC-1 IP Address>
```

5. Select **Tools > Expert Mode** and **Tools > Configuration Mode > Initial**.

Note: **Expert Mode** enables all the properties to be visible in the Signaling Manager.

Procedure:

To configure SS7 to support Scaling operations, do the following:

Attention!

Traffic will be lost for approximately 30 seconds during the SS7 stack restart. If SS7 configuration is modified, the SS7 stack is required to restart to take effect the change.



1. Back up SS7 configuration files on SC-1.

```
#cp /opt/sign/etc/active.om.cim /opt/sign/etc/active.  
om.cim.precheck.bak
```

2. LDE MIP address is used as Common Parts manager address. Refer to the following table that complies with the rule:

Navigation Pane	Operation Pane Properties	Value
System Components > System Components	CP Manager Address	ss7cafcpmaddress:6669
System Components > System Components > CP > CP	If Alias	On
System Components > System Components > ECM > ECM	Connection Time Wait	25
System Components > System Components > ECM > ECM > Services>Network Control	Pre script	/opt/sign/instance/cpm_mip_activation.sh --activate
System Components > System Components > ECM > ECM > Services>Network Control	Post script	/opt/sign/instance/cpm_mip_activation.sh --deactivate

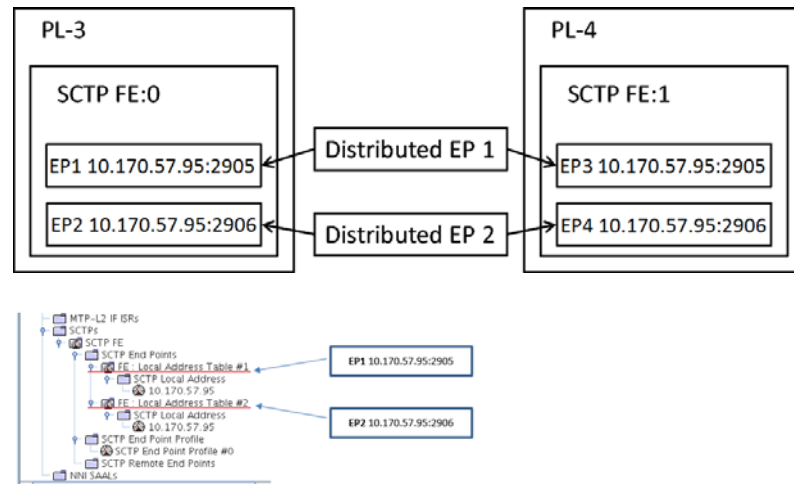
3. Asynchronous connection is enabled for BE, FE and NMP processes. Refer to the following table that complies with the rule:

Navigation Pane	Operation Pane Properties	Value
System Components > System Components > CP > CP	Msg Conn Time Wait	25
System Components > System Components > ECM > ECM> Process Classes > SCTP FEP	Command	"-w 5" should be added to launching command. For example: /opt/sign/EABss7052/bin /fe_sctp -e 255 -u 161 -a 1 -o 5 -w 5
System Components > System Components > ECM > ECM> Process Classes > GEN RP	Command	"-w 5" should be added to launching command. For example: /opt/sign/EABss7053/bin /be -b 3 -u 161 -a 5 -o 1 -d 0 -w 5
System Components > System Components > ECM > ECM> Process Classes > NMP	Command	"-w 5" should be added to launching command. For example: /opt/sign/EABss7053/bin /be -b 2 -e 255 -u 161 -a 1 -w 5

4. Only SCTP Distributed End Points (EPs) are configured in SCTP configuration.

Note: For Diameter over SCTP scenario, skip the step.

As the below SCTP Distributed End Point configuration (recommended by IPWorks) shows, two Distributed EPs are configured for SCTP FE. These two Distributed EPs use the same eVIP (e.g., 10.170.57.95) and different port (2905 and 2906).



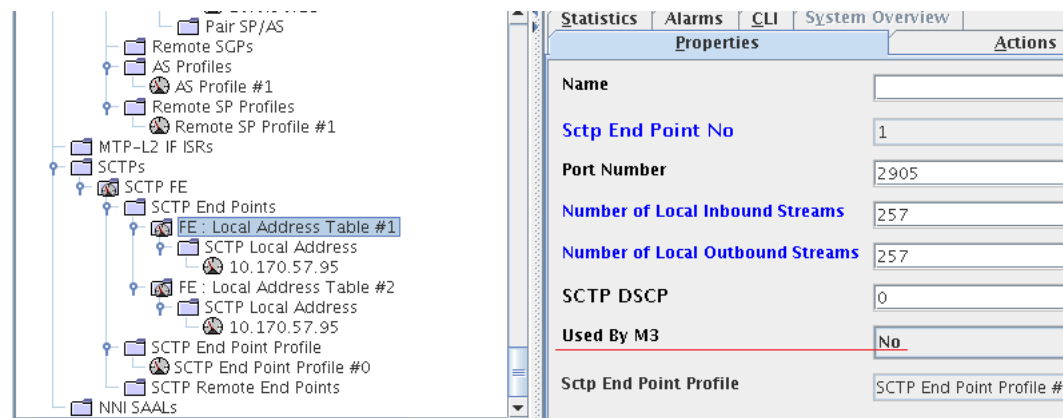
For SCTP Distributed End Points, check if below configurations are satisfied:

- a. Enable usage of eVIP functionality in SS7 CAF and usage of SCTP Distributed End Points feature.

Navigation Pane	Operation Pane Properties	Value
System Components > System Components > CP > CP	EVIP	on
M3UA IETF > M3UA	Distributed End Point Support	on

- b. For all SCTP End Points, the option **Used By M3** option is set to **No**.

For example:



For more information about the configurations of Distributed End Points, you can check the IPWorks SS7 configuration templates . Perform one of the following procedures depending on the IPWorks service to be used:

- For ENUM NP, refer to the section *Starting From Template Configuration File* in *Configure SS7 for ENUM Number Portability*.
- For AAA, refer to the section *Starting From Template Configuration File* in *Configure SS7 for AAA*.

5. If the SS7 configuration is modified, in order to make the change take effect, need to validate and restart the SS7 stack, do the following:
 - a. Validate the configuration by selecting **Edit > Validate**.
 - b. If there are validation errors, click **Results** to view error description and go to the respective configuration.
 - c. Select **Tools > Process View... > Configure** in the process view dialog box, and select **Initial Configuration** to make any update take effect.
 - d. Restart SS7 Stack on SC-1.

First, restart one SS7 stack:

```
amf-adm restart safSu=SC-1,safSg=2N,safApp=ERIC-s  
s7caf.mgmt
```

```
amf-adm restart safSu=SC-2,safSg=2N,safApp=ERIC-s  
s7caf.mgmt
```

```
amf-adm restart safSu=PL-3,safSg=2N,safApp=ERIC-s  
s7caf.netwcontrol
```



```
amf-adm restart safSu=PL-3,safSg=NWA,safApp=ERIC-ss7caf.core
```

After 1 minute, restart other ss7 stack:

```
amf-adm restart safSu=PL-4,safSg=2N,safApp=ERIC-ss7caf.netwcontrol
```

```
amf-adm restart safSu=PL-4,safSg=NWA,safApp=ERIC-ss7caf.core
```

- e. Select **File > Connect** and make sure that the status is **Active** in the status bar.
- f. Save the configuration file as another name by selecting **File> Save As**.
- g. Verify stack configuration.

- If IPWorks AAA is deployed:

Refer to the *Verify Stack Configuration* section in *Configure SS7 for AAA*.

- If IPWorks ENUM is deployed:

Refer to the *Verify Stack Configuration* section in *Configure SS7 for ENUM Number Portability*.

If verification of stack configuration fails, and the issue cannot be fixed. You need to restore the SS7 configuration with the following step:

Close the Signaling Manager, and then restore the SS7 configuration by using the SS7 configuration backup file.

```
#cp /opt/sign/etc/active.om.cim.precheck.bak /opt/sign/etc/active.om.cim
```

Repeat the steps from Step d to Step g.

1.2.2 Configuring eVIP to Support Scaling Operations

Check the eVIP configuration by using ECLI, if only EvipNode=3 and EvipNode=4 are configured under *EvipCluster*, you need to configure eVIP in ECLI.

View the information of *EvipCluster* configuration, for example:

```
# /opt/com/bin/cliss
```

```
>ManagedElement=<Node Name>,Transport=1,Evip=1,EvipDeclarations=1,EvipCluster=1
```



```
(EvipCluster=1) >show
```

```
EvipCluster=1
  commandsForAllUndesignated
    "4:set_local_port_range"
    "3:set_default_route_ipv6_sig"
    "2:set_default_route_ipv4_sig"
    "1:flush_ipv6_default"
    "0:flush_route_cache"
  primaryInterface="eth0"
EvipNode=4
EvipNode=3
```

To further configure eVIP, do the following:

1. Log on to server which launches SC-1 and then stop SC-1.

```
#ssh root@<Host1_IP_ADDRESS>
#virsh shutdown <instance name of SC-1>
```

2. When SC-1 is shut down, log on to SC-2, and add an eVIP node for SC-1.

```
SC-2:~# /opt/com/bin/cliss

> ManagedElement=<Node Name>,Transport=1,Evip=1,EvipDec
larations=1,EvipCluster=1

(EvipCluster=1) >configure

(config-EvipCluster=1) >EvipNode=1

(config-EvipNode=1) >hostname="SC-1"

(config-EvipNode=1) >commit

(config-EvipNode=1) >exit
```

3. Start the SC-1 on Host1.

```
#virsh start <instance name of SC-1>
```

4. Wait until services on SC-1 are up. MySQL Data Node must be started.

To see the MySQL cluster status:

```
# /etc/init.d/ipworks.mysql show-status
```

5. Stop the SC-2 in Host2.

```
$ ssh root@<Host2_IP_ADDRESS>
```



```
#virsh shutdown <instance name of SC-2>
```

6. When SC-2 is shut down, log on to SC-1, validate the file `evip.xml`.

```
SC-1:~ #xmllint --schema /opt/vip/etc/evipconf.xsd
/cluster/storage/system/config/evip-apr9010467/evip.xml
```

Note: If the validation is not passed, export eVIP configuration from eVIP CLI, it will overwrite the corrupted `evip.xml` file.

```
SC-1:~ # telnet `/opt/vip/bin/getactiv
econtrol` 25190
EVIP> enable
OK
EVIP# save-config
OK
EVIP# exit
```

Then, validate the file again:

```
SC-1:~ # xmllint --schema /opt/vip/etc/evipconf.x
sd /cluster/storage/system/config/evip-apr901046
7/evip.xml
```

7. When SC-2 is shut down, log on to SC-1, and add an eVIP node for SC-2.

```
SC-1:~# /opt/com/bin/cliss

> ManagedElement=<Node Name>,Transport=1,Evip=1,EvipDec
larations=1,EvipCluster=1

(EvipCluster=1)>configure

(config-EvipCluster=1)>EvipNode=2

(config-EvipNode=2)>hostname="SC-2"

(config-EvipNode=2)>commit

(config-EvipNode=2)>exit
```

8. Start the SC-2 on Host2 and wait until services on SC-2 are up.

```
$virsh start <instance-id or instance-name of SC-2>
```

9. Add other eVIP configuration on SC-1 or SC-2 by executing script.

```
#/opt/ipworks/common/scripts/add_evip_configuration.py

Add evip node ...
Add evip lbe for EvipAlb=ipw_sig_sp
Add evip se for EvipAlb=ipw_sig_sp
```



```
Add evip fee for EvipAlb=ipw_sig_sp
Add evip lbe for EvipAlb=ipw_data_sp
Add evip se for EvipAlb=ipw_data_sp
Add evip fee for EvipAlb=ipw_data_sp
Done
```

1.2.3 Checking the Tools

1.2.3.1 Check the File pxeboot.qcow2

Check if the file pxeboot.qcow2 is existed in the directory /root/auto_deployment/images. If no, do the following:

1. Create a new folder /root/auto_deployment/.

```
# mkdir -p /root/auto_deployment/
```

2. Transfer IPWorks delivered package to Host1 (for example, /root), and check md5sum of them.

```
# md5sum /root/19010-CXP9023809_3_Ux_<Revision
Number>.tar.gz
```

For example,

```
#md5sum 19010-CXP9023809_3_Ux_T.tar.gz
```

```
c296950fb749a1c863807f4945757c6e 19010-CXP9023809_3_
Ux_A.tar.gz
```

3. Unzip the package into /root/auto_deployment/ to get the qcow2 image.

```
# cd /root/auto_deployment/
```

```
# tar -zxvf /root/19010-CXP9023809_3_Ux_<Revision
Number>.tar.gz
```

For example,

```
# tar -zxvf /root/19010-CXP9023809_3_Ux_A.tar.gz
```

```
images/
images/pxeboot.qcow2
images/ipw-sc-22.qcow2
temp/
temp/mode22/
temp/mode22/ipw-vnf-22-zone.yaml
temp/mode22/ipw-vnf-22.yaml
```



1.2.3.2 Check Auto-Deploy Tool

Check if the Auto-Deploy Tool is existed. If no, do the following:

1. Create a new folder `/root/auto_deployment/<VNF_Name>`.

```
# mkdir -p /root/auto_deployment/<VNF_Name>
```

For example:

```
# mkdir -p /root/auto_deployment/IPW2
```

2. Transfer IPWorks delivered package to Host1 (for example, `/root`), and check md5sum of them.

```
# md5sum /root/19010-CXP9029034_3_Ux_<Revision>.tar.gz
```

For example,

```
# md5sum 19010-CXP9029034_3_Ux_A.tar.gz
```

```
ad013a05158f5e2931012af861ae8c77 19010-CXP9029034_3_Ux_A.tar.gz
```

3. Unzip the package to get the Auto-Deploy Tool (`ipwdeploy.sh`).

```
# cd /root/auto_deployment/IPW2
```

```
# tar -zxvf /root/19010-CXP9029034_3_Ux_<Revision>.tar.gz
```

You can find the tool in `/root/auto_deployment/IPW2/kvm_deployment/`.

1.2.4 Adapting Configuration File

This section is only applicable to the IPWorks that is upgraded from previous release version, if the IPWorks is newly deployed, skip this section.

Users need to copy the parameter values from the old file `ipwenv.conf` to the new file `ipwenv.conf`.

- Old file path: `/root/auto_deployment/kvm_deployment/config/ipwenv.conf`
- New file path: `/root/auto_deployment/<VNF_Name>/kvm_deployment/config/ipwenv.conf`

Table 1 lists the parameters used for the same purpose in the old and new `ipwenv.conf`. The newly added parameters for upgraded IPWorks are not included in this table.



Note: In the new configuration file `ipwenv.conf`, the following parameters must be empty string:

- `VNF_NAME=""`
- `SEP_CHAR=""`

Table 1 Parameter Mapping Table

Parameters in old ipwenv.conf file	Parameters in new ipwenv.conf file
PL_NUM=2	TOPO_TYPE=22
PL_NUM=4	TOPO_TYPE=222
HOST1	Host IP address in DHOST1_INFO
HOST2	Host IP address in DHOST2_INFO
HOST3	Host IP address in DHOST3_INFO
HOST4	Host IP address in DHOST4_INFO
NTP_SERVER0	NTP_SERVER0
TIME_ZONE	TIME_ZONE
IPW_INT_SP_VID	IPW_INT_SP_VID
IPW_OM_SP1_VID	IPW_OM_SP1_VID
IPW_OM_SP1_SC1_IP	IPW_OM_SP1_SC1_IP
IPW_OM_SP1_SC2_IP	IPW_OM_SP1_SC2_IP
IPW_OM_SP1_NW	IPW_OM_SP1_NW
IPW_OM_SP1_VRRP_GW_IP	IPW_OM_SP1_VRRP_GW_IP
IPW_OM_SP2_VID	IPW_OM_SP2_VID
IPW_OM_SP2_SC1_IP	IPW_OM_SP2_SC1_IP
IPW_OM_SP2_SC2_IP	IPW_OM_SP2_SC2_IP
IPW_OM_SP2_NW	IPW_OM_SP2_NW
MIP_OAM_IP	MIP_OAM_IP
MIP_PROV_IP	MIP_PROV_IP
IPW_SIG_SP1_VID	IPW_SIG_SP1_VID
IPW_SIG_SP1_NETMASK	IPW_SIG_SP1_NETMASK
IPW_SIG_SP1_PL3_IP	IPW_SIG_SP1_FEE1_IP
IPW_SIG_SP1_PL4_IP	IPW_SIG_SP1_FEE2_IP
IPW_SIG_SP1_PL5_IP	IPW_SIG_SP1_FEE3_IP
IPW_SIG_SP1_PL6_IP	IPW_SIG_SP1_FEE4_IP
IPW_DATA_SP1_VID	IPW_DATA_SP1_VID
IPW_DATA_SP1_NETMASK	IPW_DATA_SP1_NETMASK
IPW_DATA_SP1_PL3_IP	IPW_DATA_SP1_FEE1_IP
IPW_DATA_SP1_PL4_IP	IPW_DATA_SP1_FEE2_IP
IPW_DATA_SP1_PL5_IP	IPW_DATA_SP1_FEE3_IP



Parameters in old ipwenv.conf file	Parameters in new ipwenv.conf file
IPW_DATA_SP1_PL6_IP	IPW_DATA_SP1_FEE4_IP
VIP_TRF_IP1	VIP_TRF_IP1
VIP_TRF_IP2	VIP_TRF_IP2
VIP_SS7_IP1	VIP_SS7_IP1
VIP_SS7_IP2	VIP_SS7_IP2
VIP_DATA_IP	VIP_DATA_IP

1.2.5

Parameters for Scaling

Table 2 lists the parameters that are used for scaling operation.

Table 2 Parameters for Scaling Operation (2+2)

Parameter	Description	Type
QCOW2_DIR	<p>The directory for IPWork VNF packages.</p> <p>BASE_QCOW2_DIR=/root/auto_deployment/images/<VNF_NAME></p> <p>RUN_QCOW2_DIR=/root/auto_deployment/images/<VNF_NAME>/run</p> <p>Note: DO NOT mix up the two directories during the scaling. Otherwise, some unexpected issues might occur.</p>	String
TIME_ZONE	<p>In Linux OS, go to folder /usr/share/zoneinfo/ to find supported time zones: you can find Shanghai file under /usr/share/zoneinfo/Asia, more timezone info can find here: https://en.wikipedia.org/wiki/List_of_tz_database_time_zones</p>	Time Zone
NTP_SERVER0	NTP Server IP Address.	IP Address
VNF_NAME	<p>Use VNF_NAME to distinguish different installation images.</p> <p>Note: If the IPWorks is upgraded from the IPWorks (prior to 1.9) to current IPWorks release, the value of VNF_NAME must keep empty string.</p>	String
SEP_CHAR	<p>Use to separate VNF name and SC/PL information in VM name. Default value is that SEP_CHAR=".". </p> <p>Note: If the IPWorks is upgraded from the IPWorks (prior to 1.9) to current IPWorks release, the value of SEP_CHAR must keep empty string, SEP_CHAR="".</p>	String



Parameter	Description	Type
GET_PWD_FROM_CLI	<p>The default value is <code>true</code>. You only need to provide the placeholder <code>SHOST_INFO</code>.</p> <p>When the value is <code>false</code>, you should provide the true password. An example of those cases as following:</p> <ul style="list-style-type: none">When the value is <code>true</code>, you need to input root password from CLI screen. # <code>GET_PWD_FROM_CLI = "true"</code> <code>SHOST_NUM=2</code> <code>SHOST1_INFO=(10.175.161.143 fake 0)</code> <code>SHOST2_INFO=(10.175.161.144 fake 1)</code>When the value is <code>false</code>, you can set root password in <code>SHOST_INFO</code> .# <code>GET_PWD_FROM_CLI = "false"</code> <code>SHOST1_INFO=(10.175.161.143 root000 0)</code>	String
SHOST_NUM	<p><code>SHOST_NUM</code> means user will use <code>SHOST_INFO</code> to scale new PL. The default value is <code>SHOST_NUM=0</code>.</p> <p>When <code>SHOST_NUM</code> is 1, you use the <code>SHOST1_INFO</code> to scale new PL.</p> <p>When <code>SHOST_NUM</code> is 2, you use the <code>SHOST1_INFO</code> to scale new PL firstly, then use the <code>SHOST2_INFO</code> to scale new PL.</p> <p>The other <code>SHOST</code> will not be touched.</p>	Number
SHOST1_INFO SHOST2_INFO	<p>The relevant information of scaling Host. As following,</p> <p>When the <code>SHOST_NUM</code> is 1, you will use <code>SHOST1_INFO</code> to scale new PL. The third parameters 0 means using <code>Scale0</code> as the scaling VM name and <code>Scale0</code> is corresponding to PL-5.</p> <p>Accordingly, the parameter of 1 means using <code>Scale1</code> as the scaling VM name and <code>scale1</code> is corresponding to PL-6. PL number will increase orderly.</p> <p><code>SHOST1_INFO=(10.175.161.143 fake 0)</code> <code>SHOST2_INFO=(10.175.161.144 fake 1)</code></p>	Combine
PL_MEMORY	<p>PL memory size for Scale PL. Unit is GiB.</p> <p>The default value is that <code>PL_MEMORY=8</code>.</p> <p>This value must be the same with the one used for deployment.</p>	Number
PLx_VCPU_NUM	<p>PL VCPU number which must match the VCPU number calculated by Scale PL CPU set. The using VCPU number of <code>PLx_VCPU_SET</code> must be equal to <code>PLx_VCPU_NUM</code>. The <code>PLx</code> means the PL blade, such as PL5 or PL6.</p> <p>For example, When <code>PLx_VCPU_NUM=14</code>, the <code>PL5_CPU_SET=9-15,25-31</code> and <code>PL6_CPU_SET=9-15,25-31</code> or PL7, PL-8 must use 14 VCPU independently.</p> <p>The default value is <code>PLx_VCPU_NUM=14</code>.</p> <p>This value must be the same with the one used for deployment.</p>	Number



Parameter	Description	Type
SCALE0_CPU_SET SCALE1_CPU_SET	<p>SCALE0 or SCALE1 means the scale name and it is corresponding to the third parameter in SHOST_INFO.</p> <p>0 means the scale name is SCALE0. 1 means the scale name is SCALE1.</p> <p>Refer to PLx_VCPU_NUM to configure those parameters. The default value are as below:</p> <p>SCALE0_CPU_SET=1-7,17-23</p> <p>SCALE1_CPU_SET=1-7,17-23</p> <p>Note: The scale VM CPU set value depends on which NUMA node used, and would be changed.</p>	Number
MAC_ADDR_PREFIX	<p>The prefix of the MAC address. Every SC/PL blade occupy three prefix and the label starts with the number 1 as SC-1, number 2 as SC-2, number 3 as PL-3, number 4 as PL-4 and so on.</p> <p>The default value is MAC_ADDR_PREFIX="02:10:20:01:0".</p> <p>When MAC_ADDR_PREFIX=02:10:20:01:0, it is equal to following configuration:</p> <p>SC-1: 02:10:20:01:00 02:10:20:01:01 02:10:20:01:02</p> <p>SC-2 02:10:20:02:00 02:10:20:02:01 02:10:20:02:02</p> <p>PL-3 02:10:20:03:00 02:10:20:03:01 02:10:20:03:02</p> <p>PL-4 02:10:20:04:00 02:10:20:04:01 02:10:20:04:02</p>	Number
SCALE_MAX_NUM	The max number of scaling. Keep the number as 8.	Number

In addition, users need configure the parameter of PMD_CPU_MASK_PRIME and NO_PMD_CPU_MASK_PRIME in internal.conf file:

```
PMD_CPU_MASK_PRIME=0x10000
```

```
NO_PMD_CPU_MASK_PRIME=0x1
```

1.3 Related Information

Trademark information, typographic conventions, definitions, and explanations of acronyms and terminology can be found in the following documents:

- *Glossary of Terms and Acronyms*
- *Trademark Information*



- *Typographic Conventions*



2 Scaling Procedure

Following topics are included in this section:

- Section 2.1 Creating Backup before Scaling on page 17
- Section 2.2 Scale-Out Operation on page 17
- Section 2.3 Scale-In Operation on page 20
- Section 2.4 Scaling Health Check on page 25
- Section 2.5 Creating the Final Backup on page 27

Create a compute node in the IPWorks VNF as part of the scale-out operation is out of the scope of this document. However, the backup creation before and after scaling operation is required and it is a part of the scaling procedure.

For how to get the deployment and scaling tool, refer to the section *Software* in *IPWorks Auto Deployment Guideline for KVM - DL380 Gen9*.

2.1 Creating Backup before Scaling

Both System data backup and user data backup must be executed before scaling. For more details, refer to the documents *Backup and Restore*.

2.2 Scale-Out Operation

2.2.1 Overview

Scale-out means that the new VMs are instantiated. Those instances are added to the cluster automatically. Follow the instructions given by the KVM management system about how to create a VM instance.

There are two different IPWorks deployments for scaling:

- For the IPWorks before 1.9:

After upgrading to current IPWorks successfully, users must take notice of following configuration.

- In the configuration file of `ipwenv.conf`, the `VNF_NAME` must keep empty string (`VNF_NAME=""`).
- In the configuration file of `ipwenv.conf`, the `SEP_CHAR` must keep empty string (`SEP_CHAR=""`).



For more configuration details, refer to *IPWorks Auto Deployment Guideline for KVM - DL380 Gen9* to configure `VNF_NAME`

- For IPWorks 1.9 (or higher) which already supports scaling:

The length of `VNF_NAME` must be not more than five characters. There is no restriction for `SEP_CHAR`. Follow Section 2.2.2 Operating Scale-Out on page 18 to execute scaling operation.

The scale-out operation is triggered automatically once the new resource is available and launched. Once this part is covered, continue in Section 2.2.3 Monitoring the Scale-Out Progress on page 19 to monitor the state of the scale-out procedure.

2.2.2 Operating Scale-Out

Note: Scale-out operation cannot be performed if there is an existed configuration in OVS on the HOST where instances are scaled.

Execute this command to check:

```
#ovs-vsctl show
```

If existed, remove the existed configuration in OVS before scale-out.

```
#ovs-vsctl del-br <Bridge name>
```

Execute the following command to perform the scale-out operation:

1. Modify the configuration in `ipwenv.conf` first. For detail, refer to the table *IPWorks VNF Deployment Parameter List 2* in *IPWorks Auto Deployment Guideline for KVM - DL380 Gen9*.

2. Go the directory of `/root/auto_deployment/<VNF_Name>/kvm_deployment/`.

```
#cd /root/auto_deployment/<VNF_Name>/kvm_deployment/
```

For example:

```
#cd /root/auto_deployment/IPW2/kvm_deployment/
```

3. Generate pxeboot QCOW2 for scaled instances.

```
#!/ipwdeploy.sh -a genimg -T s
```

4. Copy pxeboot images and libvirt xml to the host where instances are scaled.

```
#!/ipwdeploy.sh -a prepare -T s
```

5. Start scaled instances.



- To start all instances on all SHOSTs.

```
#./ipwdeploy.sh -a scaleout
```

- Or to start instance on specific SHOST one by one.

```
#./ipwdeploy.sh -a scaleout -l <x> -m <y>
```

For example,

```
#./ipwdeploy.sh -a scaleout -l 1 -m 0
```

-l 1 means that you select SHOST1_INFO. -m 0 means that you select Scale0 PL.

2.2.3 Monitoring the Scale-Out Progress

To monitor the scale-out progress, do the following:

1. Log on to ECLI.

```
#ssh <user>@<OAM_MIP> -p <port> -s -t cli
```

2. Navigate to the Scaling Management model information:

```
>ManagedElement=<Node name>,SystemFunctions=1,SysM=1,CrM=1
```

3. Verify that the scale-out process has started.

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

```
(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-6
  adminState=UNLOCKED
  instantiationState=INSTANTIATING
  operationalState=ENABLED
  provides="ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, Role=Default-Role"
  uses="ManagedElement=1, Equipment=1, ComputeResource=PL-6"
...
```

This example shows that `instantiationState` has changed to `INSTANTIATING` for node PL-6. It means that the scale-out has started.

4. Continue to monitor the progress until the scale-out process has ended and the added node has joined the cluster:



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

For example:

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

```
ComputeResourceRole=PL-3
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-4
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-5
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-6
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=SC-1
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=SC-2
  instantiationState=INSTANTIATED
  operationalState=ENABLED
(CrM=1) >
```

2.3 Scale-In Operation

2.3.1 Overview

Scale-in means that the existing instances are removed from the cluster and then its corresponding VMs are deleted. There are two kinds of scale-in:

- Graceful scale-in (Recommended): User removes the resource in SC ECLI first and then scale in the VM by using HOT template. The traffic is gracefully switched to active PL.

For the operation steps of graceful scale-in, refer to Section 2.3.2 Graceful Scale-In Operating on page 21

- Forceful scale-in: Users must remove the resource in ECLI to clean the dirty data if one of the following cases occurs:
 - Graceful scale-in fails.
 - The VM is down or removed accidentally.



For the operation steps of forceful scale-in, refer to Section 2.3.3 Forceful Scale-In Operating on page 21

Follow the instructions given by the KVM management system about how to remove a VM instance. The scale-in operation is triggered automatically once the resource is stopped and removed.

Note:

- Scale-in only can be operated on the PLs which have been scaled out before.
- For scale-in operation, any of the existing PL in the cluster could be removed except PL3 and PL4. There is no rule to apply in this sense.

2.3.2 Graceful Scale-In Operating

To perform the graceful scale-in, do the following process:

To do the scale-in, do the following:

1. Section 2.3.4 Remove PL from Cluster on page 22

If the PL is failed to be removed due to some unexpected reason, go to the Section 2.3.3 Forceful Scale-In Operating on page 21 directly.

2. Section 2.3.5 Remove VM Instance on page 22
3. Section 2.3.6 Monitoring the Scale-In Progress on page 23

2.3.3 Forceful Scale-In Operating

To do the scale-in, do the following:

1. If the PL to be scale-in is still up, stop IPWorks service running on this PL. Take PL-5 as example:

```
#ipw-ctr status all | grep PL-5 -A20
```

```
#ipw-ctr stop <running_services> PL-5
```

For example:

```
#ipw-ctr stop aaa_radius_stack PL-5
```

2. Section 2.3.5 Remove VM Instance on page 22
3. Section 2.3.4 Remove PL from Cluster on page 22
4. Section 2.3.6 Monitoring the Scale-In Progress on page 23



2.3.4 Remove PL from Cluster

1. Log on to the ECLI.

```
#ssh <user>@<OAM_MIP> -p <port> -s -t cli
```

2. Switch to configuration mode.

```
> configure
```

3. Go to compute ComputeResourceRole MO of the PL to be scaled in.

```
(config)>ManagedElement=<Node name>,SystemFunctions=1,SysM=1,CrM=1,ComputeResourceRole=PL-<N>
```

For example:

```
(config)> ManagedElement=1, SystemFunctions=1, SysM=1, CrM=1, ComputeResourceRole=PL-5
```

4. Request to remove the PL.

```
(config-ComputeResourceRole=PL-<N>) > no provides
```

For example:

```
(config-ComputeResourceRole=PL-5) > no provides
```

5. Commit the change.

```
(config-ComputeResourceRole=PL-<N>) > up
```

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

For example:

```
(config-ComputeResourceRole=PL-5) > up
```

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

6. Monitor if the PL is removed.

Follow steps in Section 2.3.6 Monitoring the Scale-In Progress on page 23 to verify the transient states the PL goes through until this is removed.

2.3.5 Remove VM Instance

1. Go the directory of /root/auto_deployment/<VNF_Name>/kvm_deployment/.

```
#cd /root/auto_deployment/<VNF_Name>/kvm_deployment/
```

For example:



```
#cd /root/auto_deployment/IPW2/kvm_deployment/
```

2. Perform scale-in.

- To scale-in all the scaled-out PLs in SHOST_INFO list.

```
#./ipwdeploy.sh -a scalein
```

Note: In addition, Scale0 is related to PL-5, Scale1 is related to PL-6, and Scale2 is related to PL-7.

- Or to scale-in the specific PLs:

For example:

```
#./ipwdeploy.sh -a scalein -l 1 -m 0
```

-l 1 means that SHOST1_INFO is selected, -m 0 means that Scale0 PL is selected.

3. After scale-in, modify the configuration in ipwenv.conf first by remove the related SHOST1_INF.

For more detail, refer to the table *IPWorks VNF Deployment Parameter List 2* in *IPWorks Auto Deployment Guideline for KVM - DL380 Gen9*.

2.3.6 Monitoring the Scale-In Progress

1. Log on to ECLI.

```
#ssh <user>@<OAM_MIP> -p <port> -s -t cli
```

2. Navigate to the Scaling Management model information.

```
>ManagedElement=<Node name>,SystemFunctions=1,SysM=1,CrM=1
```

3. Verify that the scaling process has started.

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

The following are the example outputs for the scale-in process:



```

(CrM=1)>show -r
CrM=1
  autoRoleAssignment=ENABLED
  ...
  ComputeResourceRole=PL-5
    adminState=SHUTTINGDOWN
    instantiationState=UNINSTANTIATING
    operationalState=ENABLED
    provides="ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,Role=Default-Role"
    uses="ManagedElement=1,Equipment=1,ComputeResource=PL-5"
  ...
(CrM=1)>show -r
CrM=1
  autoRoleAssignment=ENABLED
  ...
  ComputeResourceRole=PL-5
    adminState=LOCKED
    instantiationState=UNINSTANTIATING
    operationalState=DISABLED
    provides="ManagedElement=1,SystemFunctions=1,SysM=1,CrM=1,Role=Default-Role"
    uses="ManagedElement=1,Equipment=1,ComputeResource=PL-5"
  ...

```

Note: This example shows that `instantiationState` has changed to `UNINSTANTIATING` for node `PL-5`. It means that the scale-in has started. The `adminState` changes first to `SHUTTINGDOWN` and then to `LOCKED` and `operationalState` changes to `DISABLED`.

4. Continue to monitor the progress.

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

The expected result:

```

(CrM=1)>show -m ComputeResourceRole -p instantiationState,operationalState
ComputeResourceRole=PL-3
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-4
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-6
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=SC-1
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=SC-2
  instantiationState=INSTANTIATED
  operationalState=ENABLED
(CrM=1)>

```

This example shows that node `PL-5` has disappeared. It means that `PL-5` is removed from the cluster.

However, if the scaling process fails, you will receive the following result:



```
(CrM=1)>show -m ComputeResourceRole -p instantiationState,operationalState
ComputeResourceRole=PL-3
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-4
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=PL-5
  instantiationState=UNINSTANTIATION_FAILED
  operationalState=ENABLED
ComputeResourceRole=PL-6
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=SC-1
  instantiationState=INSTANTIATED
  operationalState=ENABLED
ComputeResourceRole=SC-2
  instantiationState=INSTANTIATED
  operationalState=ENABLED
(CrM=1)>
```

This example shows that the value of `instantiationState` is changed to `UNINSTANTIATION_FAILED` for node PL-5. It means that the PL-5 is not removed from the cluster.

2.4 Scaling Health Check

The health check to be performed before and after a scaling operation is listed in this section. Also an entire IPWorks health check can be performed, for more information, refer to *IPWorks Manual Health Check*.

It is not recommended to proceed with the scaling operation if the result of the health check is not successful. For troubleshooting, refer to *IPWorks Troubleshooting Guideline*.

The checks described below can be executed at once running this script:

```
#ssh root@<SC OAM MIP>
#cd /opt/ipworks/common/scripts
#./ipw_scale_hc.sh
```

Check the output if the needed log file created. Followed the example output:



```

# ----- #
# Scaling Health Check started #
# ----- #
#####
# CHECK ping #
# ----- #
# PASSED: ping status is OK #
# ----- #
#####
# CHECK ss7 #
# ----- #
# PASSED: ss7 status is OK #
# ----- #
#####
# CHECK instantiationState #
# ----- #
# PASSED: instantiationState status is OK #
# ----- #
#####
# CHECK cmwstatusnode #
# ----- #
# PASSED: cmwstatusnode status is OK #
# ----- #
#####
# CHECK cmwscalingconf #
# ----- #
# PASSED: cmwscalingconf status is OK #
# ----- #
#####
# CHECK appl #
# ----- #
# PASSED: appl status is OK #
# ----- #
#####
# CHECK servicetype #
# ----- #
The IPWorks Service Type support scaling
# ----- #
# PASSED: servicetype status is OK #
# ----- #
#####
# CHECK evip #
# ----- #
# PASSED: evip status is OK #
# ----- #
# HEALTHCHECK: PASSED
# Logfile: /cluster/storage/no-backup/ipworks/scaling/scalehc_20170827_233634.log
# ----- #

```

1. Check that the state of the following system items at Core Middleware (Core MW) level is Status OK.

```
cmw-status node app csiass comp node sg si siass su
```

2. Check that all the SS7 processes are in Running state.

```
echo -e ' procp;\ndisconnect;\nexit' | /opt/sign/EABss7
050/bin/signmcli -own.conf=/cluster/storage/system/conf
ig/ss7caf-ana90137/etc/signmgr.cnf -online=yes
```

For example:



```

SS7 PROCESS STATES
cli> connect;
EXECUTED
cli> procp;
Process                               State
GEN RP:1 [PL-3]                       Running
GEN RP:2 [PL-4]                       Running
GEN RP:3 [PL-5]                       Running
SCTP FEP:0 [PL-3]                     Running
SCTP FEP:1 [PL-4]                     Running
SCTP FEP:2 [PL-5]                     Running
NMP:0 [PL-3]                          Running
OAMP:0 [PL-3]                         Running
LOGD:0 [PL-3]                         Running
ECM:0 [PL-3]                          Running
ECM:1 [PL-4]                          Running
ECM:2 [PL-5]                          Running
ECSP:0 [PL-3]                         Running
ECSP:1 [PL-4]                         Running
ECSP:2 [PL-5]                         Running
SAFOAM:0 [PL-3]                       Running
cli> disconnect;
EXECUTED
cli> exit;

```

2.5 Creating the Final Backup

Create a backup after the scaling is performed following the same steps as described in Section **Create Initial Backup**, name it `AFTER_SCALE_PL_<Numberof_PLS_after_Scaling>`.





Reference List

- [1] *IPWorks Initial Configuration*, 5/1553-AVA 901 33/3
- [2] *IPWorks Manual Health Check*
- [3] *IPWorks Troubleshooting Guideline*
- [4] *Configure SS7 for AAA*
- [5] *Configure SS7 for ENUM Number Portability*
- [6] *IPWorks Auto Deployment Guideline for KVM - DL380 Gen9*, 19/1553-AVA 901 33/3 Uen
- [7] *BRF-C Management Guide*, 9/1553-APR 901 0444/4