

Emergency Recovery Procedure for IPWorks

EMERGENCY RECOVERY

Copyright

© Ericsson 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

1	Introduction	1
1.1	Prerequisites	1
1.2	Related Information	3
2	Emergency Definitions	5
3	Scenarios for Recovery	7
3.1	Recovery Scenarios for Both CEE and KVM	9
3.2	Recovery Scenarios for CEE	14
3.3	Recovery Scenarios for KVM	16
4	Recovery Procedures	19
4.1	Node Reboots Cyclically	20
4.2	Node PXE Booting Fails	25
4.3	PL Node Does Not Start Installation from SC By PXE	25
4.4	IPWorks System Data Restore Fails	27
4.5	IPWorks User Data Restore Fails	27
4.6	IPWorks Application Component Cannot Start	30
4.7	IPWorks Application Component Patch Fails	31
4.8	Recovering Two Damaged SC VMs	32
4.9	Hard Reboot Instance	33
4.10	Blade Replacement	35
4.11	IPWorks environment file, images and yaml file missing	42
4.12	Data Node Can Not Start Up During CEE Upgrade	44
4.13	Recovering One Damaged SC VM	45
4.14	Recovering One Damaged PL VM	53
4.15	Recovering SC VM and PL VM in One Host	54
5	Problem Reporting	59
5.1	Problem Solved	59
5.2	Consult Next Level of Support	59
6	Appendix	61
	Reference List	63





1 Introduction

This document gives an overview of the emergency recovery tasks to be performed on the IPWorks that is deployed on the Ericsson Cloud Execution Environment (CEE) or KVM on HP DL380 Gen9 and Gen10 systems.

Typically, an emergency procedure is required for conditions that make communication or normal management and alarm handling impossible. In a Worst Case Scenario, a procedure is required to restore the product. Emergency in this document refers to the situations described in Section 3 on page 7.

Scope

This document focuses on the hardware, platform, and application recovery. The network level recovery is out of the scope.

The system is assumed to have been in a fully working state before the problems started. Therefore no troubleshooting procedures that relate to faulty configuration or incorrect software version or hardware version, or both, are explained. For this type of information, refer to [IPWorks Configuration Management](#) and [IPWorks Troubleshooting Guideline](#).

Some steps that have been identified as risky from an In-Service Performance (ISP) point of view are avoided in this document. When such steps are necessary, it is recommended to contact the next level of support, see Section 5.2 Consult Next Level of Support on page 59. Thus, at least two levels of support are involved before making a risky decision.

The recovery actions described in the recovery procedures are expected to be executed by Ericsson local or global support organizations, or both.

Note: Problematic situations and all the recovery actions that have been taken should be carefully documented.

Target Groups

This document is intended for telecommunication technicians authorized to perform emergency recovery procedures on Ericsson IPWorks systems.

1.1 Prerequisites

This section states the prerequisites for performing the emergency recovery procedures.

1.1.1 Personnel

The personnel performing the emergency recovery procedure must have solid knowledge of and training in the following areas:



- Telecommunication concepts, terminology, and abbreviations
- Ericsson BSP system (only for CEE)
- Ericsson Cloud Execution Environment (only for CEE)
- IPWorks system
- Linux

1.1.2 Documents

Before starting this procedure, ensure that the following information or documents are available:

- Information about, for example, product name, software version, platform, operating system, and hardware.
- For information about how to collect data and log files, refer to [Data Collection Guideline for IPWorks](#).
- For information about how to perform backup and restore procedures, refer to [Backup and Restore](#).
- Some of the recovery steps require physical access to the products for pressing buttons, replacing hardware, and so on. For information on physical access and handling, refer to [Personal Health and Safety Information and System Safety Information](#).
- Some of the recovery steps require instructions. Refer to the following documents:
 - [IPWorks Auto Health Check](#)
 - [IPWorks Manual Health Check](#)
 - [Configure MySQL NDB Cluster](#)
 - [Create Backup](#)
 - [Restore Backup](#)

1.1.3 Tools

The following tools are required:

- Active Library Explorer (ALEX) account
- Knowledge Centered Sharing (KCS)/ Primus account
- Secure Shell (SSH) client (for remote access)



- Sametime (for remote access)

In addition, verify that all network, hardware, and cables are free of faults.

1.1.4

Access

The following access information is required for both on-site and remote access:

- IP address
- The password of the IPWorks account
- The password of the root account for OS (only for CEE)
- The password of the root account for BSP (only for CEE)
- The password of the Alas administrator account (only for CEE)
- The password of the root account for Host OS (only for KVM)

Note:

- Ensure that IP connectivity between the IPWorks and the management terminal has been correctly established before attempting to perform emergency handling procedures.
- Ensure the console port connection between the IPWorks nodes and the management terminal has been correctly established before attempting to perform emergency handling procedures.
- Cloud emergency recovery, refer to Emergency Recovery Procedure.

1.2

Related Information

Definition and explanation of acronyms and terminology, trademark information, and typographic conventions can be found in the following documents:

- Glossary of Terms and Acronyms
- Trademark Information
- Typographic Conventions





2 Emergency Definitions

Emergency in this document refers to the situation when a loss of service occurs in an IPWorks application.

Refer to following documents to handle situations where some processes cannot be started or where redundancy is affected but traffic and provisioning can still continue:

- IPWorks Troubleshooting Guideline
- Backup and Restore
- Data Collection Guideline for IPWorks





3 Scenarios for Recovery

This section describes different scenarios from which the system must recover, as shown in Table 1.

Table 1 Scenarios for Recovery

Scenario	Symptom	Recovery Procedure
Node Reboots Cyclically	A node keeps rebooting cyclically.	See Section 4.1 Node Reboots Cyclically on page 20
PXE Booting the System Fails	DHCP or TFTP fails when perform PXE boot for the system.	See Section 4.2 Node PXE Booting Fails on page 25
PL Node Does Not Start Installation from SC by PXE	A PL node has successfully booted and connected to an SC node by PXE, but does not start installation. Also, the PL node keeps returning an error message.	See Section 4.3 PL Node Does Not Start Installation from SC By PXE on page 25
IPWorks System Data Restore Fails	IPWorks System Data restore fails.	See Section 4.4 IPWorks System Data Restore Fails on page 27
IPWorks User Data Restore Fails	IPWorks User Data restore fails.	See Section 4.5 IPWorks User Data Restore Fails on page 27
IPWorks Application Component Cannot Start	IPWorks application component cannot be started by using ipw-ctr.	See Section 4.6 IPWorks Application Component Cannot Start on page 30
IPWorks Application Component Patch Fails	IPWorks application component patch fails to install.	See Section 4.7 IPWorks Application Component Patch Fails on page 31
Node Damaged	IPWorks VNF Node is damaged unexpected.	See Section 4.9 Hard Reboot Instance on page 33
Blade is Damaged	BSP GEP blade is damaged unexpectedly.	See Section 4.10 on page 35
IPWorks environment file, images and yaml file missing	IPWorks environment file, images and yaml file are missing.	See Section 4.11 IPWorks environment file, images and yaml file missing on page 42



Table 1 Scenarios for Recovery

Scenario	Symptom	Recovery Procedure
One SC VM is Damaged	SC VM cannot be accessed from console (without console port hardening). And the VM cannot be accessed by SSH from a healthy VM (like SC or PL through internal network) or from external network. The command <code>tipc-config -n</code> on the healthy VM shows the node is down.	See Section 4.13 on page 45
One PL VM is Damaged	PL VM cannot be accessed by console (without console port hardening). And the VM cannot be accessed by SSH from a healthy VM (like SC or PL through internal network) or from external network. The command <code>tipc-config -n</code> on the healthy VM shows the node is down.	See Section 4.14 on page 53
SC VM and PL VM in One Host are Damaged	IPWorks SC VM and PL VM in one host cannot be accessed from console (without console port hardening). And the VMs cannot be accessed by SSH from a healthy VM (like SC or PL through internal network) or from external network. The command <code>tipc-config -n</code> on the healthy VM shows the node is down.	See Section 4.15 on page 54
Two SC VMs are Damaged	IPWorks VM cannot be accessed from console (without console port hardening). And the VM cannot be accessed by SSH from external network.	See Section 4.8 on page 32

- Recovery scenarios for both CEE and KVM, see Section 3.1 on page 8.
- Recovery scenarios for CEE, see Section 3.2 on page 14.
- Recovery scenarios for KVM, see Section 3.3 on page 16.



3.1 Recovery Scenarios for Both CEE and KVM

3.1.1 Node Reboots Cyclically

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	A node keeps rebooting cyclically.
Possible Reasons:	<p>The possible reasons are the following:</p> <ul style="list-style-type: none"> • The <code>cluster.conf</code> file contains a wrong configuration. • The <code>evip.xml</code> file contains a wrong configuration. • Some component cannot start up and then escalate to reboot node to try to recover it. (for example, COM, Storage Server, ENUN, DNS service). • The <code>/dev/drbd</code> folder is lost. • The BIOS boot order is incorrect. • No free memory left. • No disk space left.
Recovery procedures:	Section 4.1 Node Reboots Cyclically on page 20
Risks:	In case the <code>/dev/drbd</code> folder is lost, reinstallation of the whole system might be necessary.
Duration:	<ul style="list-style-type: none"> • Recover the <code>cluster.conf</code> file and then reboot system: about 30 minutes. • Reboot the system from the GRUB boot loader: about 20 minutes. • In case the <code>/dev/drbd</code> folder is lost, the duration depends on the actions that must be taken to resolve the issue.
Expected outcome:	The system boots correctly with no cyclic reboot.



3.1.2 Node PXE Booting Fails

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	<p>When try to PXE boot the system, one node cannot boot from networks.</p> <p>1. DHCP failed as below:</p> <pre>Network boot from Intel E1000 Copyright (C) 2003-2005 VMware, Inc. Copyright (C) 1997-2000 Intel Corporation CLIENT MAC ADDR: 00 0C 29 A3 D0 09 GUID: PXE-E51: No DHCP or proxyDHCP offers were PXE-M0F: Exiting Intel PXE ROM. Network boot from Intel E1000 #2 Copyright (C) 2003-2005 VMware, Inc. Copyright (C) 1997-2000 Intel Corporation CLIENT MAC ADDR: 00 0C 29 A3 D0 13 GUID: DHCP...␣</pre> <p>2. DHCP connection is successful but cannot download boot image.</p>
Recovery procedures:	Section 4.2 Node PXE Booting Fails on page 25
Risks:	Not applicable.
Duration:	About 20 minutes.
Expected outcome:	The system PXE boots successfully with no error message displayed.

3.1.3 PL Node Does Not Start Installation from SC By PXE

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable



Symptom:	A PL node has successfully booted and connected to an SC node by PXE, but does not start installation.
Possible Reasons:	<p>The possible reasons are the following:</p> <ul style="list-style-type: none"> • A symbolic boot link is on the SC node dedicated to this PL node used at booting. • The DHCP service does not start properly. • The backplane port is disabled.
Recovery procedures:	Section 4.3 PL Node Does Not Start Installation from SC By PXE on page 25
Risks:	Not applicable
Duration:	About 30 minutes.
Expected outcome:	The PL node successfully starts installation.

3.1.4 IPWorks System Data Restore Fails

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	IPWorks system data restore fails.
Possible Reasons:	<p>The possible reasons are the following:</p> <ul style="list-style-type: none"> • System backup file is broken. • System backup of an abnormal system
Recovery procedures:	Before applying this recovery, you must contact next level of support firstly.
Risks:	Restore would fail for unexpected reason.
Duration:	<p>About 30 minutes for only system restore.</p> <p>It depends on the following:</p> <ul style="list-style-type: none"> • Whether IPWorks User Data backup file includes NDB data • The data size in NDB.
Expected outcome:	The restore operation returns the "PERMIT_PHASE is completed", for details, refer to Restore Backup .



3.1.5 IPWorks User Data Restore Fails

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	IPWorks User Data restore fails. ECLI returns error message for the restore operation.
Possible Reasons:	<p>The possible reasons are the following:</p> <ul style="list-style-type: none">• IPWorks Backup file is broken.• MySQL NDB processes are not running.• Unknown system problem.
Recovery procedures:	Before applying this recovery, you must contact next level of support firstly.
Risks:	Restore would fail for unexpected reason.
Duration:	About 30 minutes.
Expected outcome:	The restore operation returns the "PERMIT_PHASE is completed", for details, refer to Restore Backup .

3.1.6 IPWorks Application Component Cannot Start

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	<p>AMF SI Unassigned alarm which is related to Storage Server, DNS, ENUM, and so on.</p> <p>For more information on checking active alarms in the system, refer to Check Alarm Status.</p>
Notification/Event:	Not applicable
Symptom:	<p>Command <code>ipw-ctr start [comp] [<hostname>]</code> cannot start IPWorks application component.</p> <p>The following command does not return empty:</p> <pre>#ipw-ctr status [comp] [<hostname>] grep saAmfSUPresenceState grep FAIL</pre>
Possible Reasons:	



Recovery procedures:	Section 4.6 IPWorks Application Component Cannot Start on page 30
Risks:	Not applicable
Duration:	About 30 minutes.
Expected outcome:	The IPWorks application is working properly.

3.1.7 IPWorks Application Component Patch Fails

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	IPWorks application component patch failed and cannot roll back.
Possible Reasons:	<p>The possible reasons are the following:</p> <ul style="list-style-type: none"> • The IPWorks patch fails to install and fails to be rolled back. • The IPWorks path causes another big problem need to be removed.
Recovery procedures:	Section 4.7 IPWorks Application Component Patch Fails on page 31
Risks:	Restore would fail for unexpected reason.
Duration:	<p>About 30 minutes for only system restore.</p> <p>It depends on the following:</p> <ul style="list-style-type: none"> • Whether IPWorks User Data backup includes NDB • The data size in NDB
Expected outcome:	IPWorks is determined to be healthy, and the patched application can provide service.

3.1.8 Two SC VMs are Damaged

Hardware Platform:	BSP CEE or KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable



Symptom:	IPWorks VM cannot be accessed from console (without console port hardening). And the VM cannot be accessed by SSH from external network.
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.8 on page 32
Risks:	Not applicable
Duration:	About 2 hours
Expected outcome:	IPWorks system can normally handle traffic.

3.2 Recovery Scenarios for CEE

3.2.1 Node Damaged

Hardware Platform:	BSP CEE
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	<p>IPWorks VNF node cannot be accessed from console (without console port hardening).</p> <p>And the node cannot be accessed by SSH from healthy node (like SC or PL through internal network) or from external network. The command "tipc-config -n" on one healthy node shows the node is down.</p>
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.9 Hard Reboot Instance on page 33
Risks:	Not applicable
Duration:	About 1 hour.
Expected outcome:	The new VNF node can power on and boot up, and "tipc-config -n" can find it.

3.2.2 Blade is Damaged

Hardware Platform:	BSP CEE
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable



Notification/Event:	Not applicable
Symptom:	GEP blade cannot be accessed by console (without console port hardening) and cannot recover by reboot, power cycle.
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.10 on page 35
Risks:	Not applicable
Duration:	Not applicable
Expected outcome:	The new GEP blade replaced can power on and boot up.

3.2.3 IPWorks environment file, images and yaml file missing

Hardware Platform:	BSP CEE
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	IPWorks environment file, images and yaml file are missing.
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.11 IPWorks environment file, images and yaml file missing on page 42
Risks:	Not applicable
Duration:	Not applicable
Expected outcome:	The missing files are recovered.

3.2.4 Data Node Can Not Start Up During CEE Upgrade

Hardware Platform:	BSP CEE
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable



Symptom:	Node 28 or 27 is not connected. <pre>SC-1:~ # /etc/init.d/ipworks.mysql show-status Connected to Management Server at: localhost:1186 Cluster Configuration ----- [ndbd(NDB)] 2 node(s) id=27 @169.254.100.1 (mysql-5.6.31 ndb-7.4.12, No id=28 (not connected, accepting connect from SC-2) [ndb_mgmd(MGM)] 2 node(s) id=1 @169.254.100.1 (mysql-5.6.31 ndb-7.4.12) id=2 @169.254.100.2 (mysql-5.6.31 ndb-7.4.12) [mysqld(API)] 24 node(s) id=3 @169.254.100.1 (mysql-5.6.31 ndb-7.4.12) id=4 (not connected, accepting connect from SC-2)</pre>
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.12 on page 44
Risks:	Not applicable
Duration:	20 minutes
Expected outcome:	Data node starts up normally after CEE upgrade.

3.3 Recovery Scenarios for KVM

3.3.1 One SC VM is Damaged

Hardware Platform:	KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	SC VM cannot be accessed from console (without console port hardening). And the VM cannot be accessed by SSH from a healthy VM (like SC or PL through internal network) or from external network. The command <code>tipc-config -n</code> on the healthy VM shows the node is down.
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.13 on page 45



Risks:	Not applicable
Duration:	About 1 hour
Expected outcome:	The new SC VM node can power on and boot up, and can be found by command <code>tipc-config -n</code> .

3.3.2 One PL VM is Damaged

Hardware Platform:	KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	PL VM cannot be accessed by console (without console port hardening). And the VM cannot be accessed by SSH from a healthy VM (like SC or PL through internal network) or from external network. The command <code>tipc-config -n</code> on the healthy VM shows the node is down.
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.14 on page 53
Risks:	Not applicable
Duration:	About 30 minutes
Expected outcome:	The new PL VM can power on and boot up, and can be found by command <code>tipc-config -n</code> .

3.3.3 SC VM and PL VM in One Host are Damaged

Hardware Platform:	KVM on HP DL380 Gen9 / Gen10
Operating System:	SUSE Linux Enterprise Server 12 (x86_64)
Alarm:	Not applicable
Notification/Event:	Not applicable
Symptom:	SC VM and PL VM in one host cannot be accessed from console (without console port hardening). And the VMs cannot be accessed by SSH from a healthy VM (like SC or PL through internal network) or from external network. The command <code>tipc-config -n</code> on the healthy VM shows the node is down.
Possible Reasons:	Not applicable
Recovery procedures:	Section 4.15 on page 54



Risks:	Not applicable
Duration:	About 1.5 hours
Expected outcome:	The new PL VM and SC VM can power on and boot up, and can be found by command <code>tipc-config -n</code> .

3.3.4 Two SC VMs are Damaged

Refer to Section 3.1.8 on page 13.



4 Recovery Procedures

The procedures in this section describe the various scenarios used to find and resolve faults that can cause an IPWorks emergency situation.

The execution of the emergency recovery procedure follows the workflow as described in the following steps and shown in Figure 1:

1. Identify the problem type best matching the problem experienced.
2. Identify the recovery scenario best matching the problem experienced.
3. Execute recovery actions in increasing order of severity.
4. If recovery is successful, take preventive actions to prevent the problem from reoccurring.




Figure 1 Workflow

4.1 Node Reboots Cyclically

Perform the recovery procedure according to the following scenario:

1. In case of `cluster.conf` has invalid configuration, check current `cluster.conf` configuration with template and make sure all network info is correct. Then, you can modify the `cluster.conf` in maintenance mode, see Section 4.1.2 Rebooting the System from GRUB Boot Loader on page 23.
2. In case of `evip.xml` has invalid configuration. Modify the `/cluster/storage/system/config/evip-apr9010467/evip.xml` by following the IP plan and then validate the schema by executing command `xmllint -schema /opt/vip/etc/evipconf.xsd /cluster/storage/system/config/evip-apr9010467/evip.xml`. It will



print whole `evip.xml` body if no error found. This issue only occurs in PL because SC does not use eVIP function.

3. In case application keeps restarting and escalating to node reboot, follow the steps below:
 - a. Investigate which application causes the node reboot.
 - b. Stop the application in the node which is not cyclically reboot, see Section 4.1.1 Stop Application Causing Node Cyclic Reboot on page 21.
4. In case the `/dev/drbd` folder is lost, check the DRBD configurations by `cat /proc/drbd`. Execute command "`drbd-overview`" to get drbd overview status. User can perform power cycle to try to recover it. See Section 4.9 Hard Reboot Instance on page 33 for details. If the problem is not resolved, see Section 5.2 Consult Next Level of Support on page 59. Reinstallation and IPWorks restoration might be necessary.
5. In case of no free memory left, after system reboot, check system memory usage by using `top` command to find whether this process occupies much memory, and then stop it. Refer to Section Checking CPU and Memory in IPWorks Manual Health Check.
6. In case of no free disk left, after system reboot, check the system disk usage in rebooting node console terminal. Refer to the Section Checking Disk Usage in IPWorks Manual Health Check.
7. In all other cases, go to Section 4.1.2 Rebooting the System from GRUB Boot Loader on page 23. And then go to OS maintenance mode to find what is wrong, like disk usage, disk label and configuration for `cluster.conf`.

4.1.1 Stop Application Causing Node Cyclic Reboot

It is possible that IPWorks application failed to restart and then escalate to node reboot to try to recover the system automatically. But according to that the failure still exists, so this causes the node reboots again and again. When node cyclic reboot occurs, check the node-related log to find the possible cause first.

1. Check console terminal output.

Find if there is any abnormal or any fail, error message in console terminal output. Show status of IPWorks application status, such as Storage Server, DNS, ENUM, DNS SM, AAA SM, in SC or PL.

2. Check the reboot node messages log.

SC messages log is `/var/log/<SC-ID>/messages`, if only one SC cyclic reboot, check the log content in another healthy SC because the message can be accessed in both SC. PL messages log can also be accessed in both SC nodes in `/var/log/<PL-ID>/messages`.

Note: `<SC-ID>` can be SC-1 or SC-2, `<PL-ID>` can be PL-3 or PL-4 etc.



3. Search failed application information.

Search `restart`, `reboot`, `recovery`, `escalate` in messages log to find which IPWorks component triggers node reboot.

4. Stop application which causes cyclic reboot.

The `ipw-ctr` command can be executed in any node of IPWorks cluster nodes (both SC and both PL). So, execute the command to stop application in the healthy node. Or this command can be executed in console terminate.

```
#ipw-ctr stop <IPW application> <hostname>
```

For example:

— Stop Storage Server in SC-1:

```
#ipw-ctr stop ss SC-1
```

— Stop ENUM/DNS in PL-3:

```
#ipw-ctr stop enum PL-3;ipw-ctr stop dns PL-3
```

— Stop AAA Diameter in PL-3:

```
#ipw-ctr stop aaa_diameter PL-3
```

5. Troubleshoot why the application fails to start by checking application log.

For Storage Server in SC:

Refer to the Section Failed to Stop/Start/Restart Storage Server by `ipw-ctr` in IPWorks Troubleshooting Guideline.

For DNS Server in PL:

```
>dn ManagedElement=1,IpworksFunction=1,IpworksDnsRoot=1
,DnsServer=1,BindService=1,DnsLog=1
>configure
(config-DnsLog=1)>level=DNS_LOG_LEVEL_DEBUG
```

The log can be found in `/cluster/storage/no-backup/ipworks/logs/<PL-ID>`.

Refer to the Section DNS Server Fails to Start after System Boot in IPWorks Troubleshooting Guideline.

For ENUM Server in PL:

Modify log level in ECLI:

```
> dn ManagedElement=1,IpworksFunction=1,IpworksDnsRoot=
1,IpworksEnumRoot=1,EnumServer=1,Log=1
```



```
>configure
(config-Log=1)>level=LOG_LEVEL_TRACE
```

Refer to the Section Failed to Stop/Start/Restart ENUM Server by ipw-ctr in IPWorks Troubleshooting Guideline.

For AAA Server in PL:

Modify log level in ECLI (use PL-3 as example):

```
>ManagedElement=1,IpworksFunction=1,IPWorksAAARoot=1,IPWorksAAA
CommonRoot=1,AAAServer=PL-3,LogManagement=1,IPWorksLog=AAA_DI
AMETER_SERVER
```

```
(IPWorksLog=AAA_DIAMETER_SERVER)>configure
```

```
(config-IPWorksLog=AAA_DIAMETER_SERVER)> level=LOG_LEVEL_DEBUG
```

The log can be found in /cluster/storage/no-backup/ipworks/logs/<PL-X>.

Refer to Section AAA Server in IPWorks Troubleshooting Guideline.

4.1.2 Rebooting the System from GRUB Boot Loader

Rebooting the system from the GRUB boot loader is necessary in the following scenarios:

- A Node reboots cyclically (see Section 3.1.1 on page 9)
- PXE booting the system fails (see Section 3.1.2 on page 10)

To reboot the system from the GRUB boot loader, follow the steps below:

1. Go to the GRUB boot loader and initiate booting the system.
2. Press any key when the following text is displayed:

Press any key to continue.
3. When the GRUB boot menu is displayed, select **Maintenance mode (Serial console)**.

Note: To see the GRUB boot menu, a serial console must be attached to the machine.



```
GNU GRUB version 2.02~beta2

+-----+
| SUSE Linux Enterprise Server - Operational mode (Serial console) |
| SUSE Linux Enterprise Server - Operational mode (VGA console)   |
| *SUSE Linux Enterprise Server - Maintenance mode (Serial console)|
| SUSE Linux Enterprise Server - Maintenance mode (VGA console)   |
+-----+

Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the commands
before booting or 'c' for a command-line.
```

Note: Do not select **Maintenance mode (VGA console)** which will be hung.

4. Log on to the system as root in maintenance mode.

The password for logging in as root is rootroot.

```
[ OK ]Started LSB: Early LDE configuration.
[ OK ]Reached target Rescue Mode.
Welcome to rescuGive root password for maintenance (or press Control-D to co
linux:~#
```

5. Execute the following command:

```
# cluster config --create-devices
```

6. Execute the following command:

```
# swapon /dev/part_swap
```

7. Mount /boot using the following command:

```
# mount -t ext3 -o data=journal,commit=1 /dev/part_boot /boot
```

8. Update the cached version of the `cluster.conf` file with the new input by entering the following:

```
# vi /boot/.cluster.conf
```

9. Reboot the system.

```
#reboot
```

Note: When the system has rebooted, update the original `cluster.conf` file in `/cluster/etc` with the new input as in Step 8.

In case the problem was that a node kept rebooting cyclically, go to Step 11. If the problem was that PXE booting the system failed, go to Step 10.



10. Initiate PXE booting for the system again.
11. When the procedure is completed, do the following:
 - In case the problem is resolved, go to Section 5.1 Problem Solved on page 59.
 - In case the problem is not resolved, go to Section 5.2 Consult Next Level of Support on page 59.

4.2 Node PXE Booting Fails

Perform the recovery procedure according to the PL node boot fails.

Check whether the `ipw_1de_sp` network configuration in BSP is correct.

4.3 PL Node Does Not Start Installation from SC By PXE

Perform the recovery procedure according to the following scenario:

- In case a symbolic boot link is missing on the SC node, go to Section 4.3.1 Creating a New Symbolic Boot Link to PL Node on page 25.
- In case the DHCP service does not start properly, go to Section 4.3.2 Restart DHCP Server on page 26.
- In case path in `ipw_int_sp` is broken between PL and SC, go to Section 4.3.3 Check IPW_INT_SP Connection on page 27

4.3.1 Creating a New Symbolic Boot Link to PL Node

To create a new symbolic boot link to a PL node that does not start installation after connecting to an SC node, follow the steps below:

1. Go to the following directory on one of the SC nodes:

```
# cd /cluster/nodes/<PL_id>/
```

The variable `<PL_id>` refers to the PL node that does not start installation after connecting to the SC node.

For example, if PL3 is the node that does not start installation after connecting to the SC node, execute the following command:

```
# cd /cluster/nodes/3/
```

2. Delete the symbolic boot link dedicated to the PL node:

```
# rm boot
```



Example boot link:

```
boot -> ../sw/linux-payload-R3B02- 0  
/fa2a5eab751fa45fe91b4417e59cab5e
```

3. Go to the directory of another PL node, for example, PL-4:

```
# cd /cluster/nodes/4/
```

4. List the boot link for this PL node (PL-4):

```
# ls -l boot
```

Example output:

```
lrwxrwxrwx 1 root root 61 Aug 23 11:46 boot ->  
../sw/linux-payload-R3B02-  
0/fa2a5eab751fa45fe91b4417e59cab5e
```

5. Go back to the directory of the PL node that does not start installation (PL-3):

```
# cd /cluster/nodes/3/
```

6. Create a new symbolic boot link using the path for the other PL node (PL-4), for example:

```
#ln -s ../sw/linux-payload-R3B02-0/fa2a5eab751fa45fe91b4417  
e59cab5e boot
```

7. When the procedure is completed, do the following:

- In case the problem is resolved, go to Section 5.1 Problem Solved on page 59.
- In case the problem is not resolved, go to Section 5.2 Consult Next Level of Support on page 59.

4.3.2 Restart DHCP Server

Execute the following command on one SC node to restart DHCP server:

```
# systemctl restart dhcpd.service
```

Execute the following command to check the DHCP status:

```
# systemctl status dhcpd.service
```



4.3.3 Check IPW_INIT_SP Connection

Check SC /var/log/messages to find if there is any DHCP and TFTP log when PL node tries to boot by using PXE, if there is not any log, the network connection in IPW_INIT_SP is broken. For this kind of issue, the L2 connection generated by cloud BSP plug-in is broken, contact the cloud administrator for support.

4.4 IPWorks System Data Restore Fails

Prerequisites:

Before the IPWorks System Data restore, periodical System Data backup and User Data backup (with or without NDB) have been performed. And several backup files are created before the restore operation.

Perform the recovery procedure according to the following scenario:

- In case System Data backup file is broken, go to Section 4.4.1 Perform System Data Backup By Selecting Another Backup File on page 27.
- In case the system backed up in System Data backup file is abnormal, go to Section 4.4.1 Perform System Data Backup By Selecting Another Backup File on page 27.
- In case system always fails to be restored and cannot be recovered, IPWorks VNF needs to be re-deployed, refer to IPWorks Deployment Guide. After IPWorks VNF is re-deployed, you need to restore both System Data backup and User Data backup (with NDB), refer to [Restore Backup](#).

4.4.1 Perform System Data Backup By Selecting Another Backup File

1. Select another system backup file and then perform the restore operation again. Refer to [Restore Backup](#) for details.

It is recommended to restore the system backup file that is generated just after the whole IPWorks installation and initial configuration.

2. Start MySQL NDB Cluster manually. Refer to [Configure MySQL NDB Cluster](#) for details.

Note: IPWorks MySQL NDB does not start automatically after IPWorks System Data restore.

4.5 IPWorks User Data Restore Fails

Prerequisites:

Before the IPWorks User Data restore, periodical System Data backup and User Data backup (with or without NDB) have been performed. And several backup files are created before the restore operation.



Perform the recovery procedure according to the following scenario:

- Check `/cluster/storage/no-backup/ipworks/logs/<SC hostname>/ipwbrf.log` on the inactive SC which does not execute `com_fuse_module` process to find the possible cause.
- In case User Data backup file is broken, perform User Data backup by selecting another backup file. Refer to [Restore Backup](#) for details.
- In case MySQL NDB processes are not running, go to Section 4.5.1 Recover MySQL NDB for IPWorks User Data Restore on page 28.
- In case User Data restore always fails and cannot be recovered, go to Section 5.2 Consult Next Level of Support on page 59.

4.5.1 Recover MySQL NDB for IPWorks User Data Restore

Recovering MySQL NDB for IPWorks User Data restore is necessary in the following scenarios:

- IPWorks User Data Restore Fails (see Section 3.1.5 on page 12)
- IPWorks Application Component Patch Fails (see Section 3.1.7 on page 13)

To start IPWorks User Data restore, all MySQL NDB nodes in both SCs must start up first. However, if you want to restore IPWorks MySQL NDB data while MySQL NDB is already crashed, do the following to recover MySQL NDB to the "running" status first:

1. Check MySQL NDB status.

```
<SC hostname>:~ # /etc/init.d/ipworks.mysql show-status

Connected to Management Server at: localhost:1186
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=27   @169.254.101.1  (mysql-5.6.31 ndb-7.4.12, Nodegroup: 0, *)
id=28   @169.254.101.2  (mysql-5.6.31 ndb-7.4.12, Nodegroup: 0, Master)

[ndb_mgmd(MGM)]  2 node(s)
id=1      @169.254.101.1  (mysql-5.6.31 ndb-7.4.12)
id=2      @169.254.101.2  (mysql-5.6.31 ndb-7.4.12)

[mysqld(API)]    24 node(s)
id=3      @169.254.101.1  (mysql-5.6.31 ndb-7.4.12)
id=4 (not connected, accepting connect from SC-2)
id=5 (not connected, accepting connect from any host)
id=6 (not connected, accepting connect from any host)
id=7 (not connected, accepting connect from any host)
id=8 (not connected, accepting connect from any host)
id=9 (not connected, accepting connect from any host)
```




```
id=10 (not connected, accepting connect from any host)
id=11 (not connected, accepting connect from any host)
id=12 (not connected, accepting connect from any host)
id=13 (not connected, accepting connect from any host)
id=14 (not connected, accepting connect from any host)
id=15 (not connected, accepting connect from any host)
id=16 (not connected, accepting connect from any host)
id=17 (not connected, accepting connect from any host)
id=18 (not connected, accepting connect from any host)
id=19 (not connected, accepting connect from any host)
id=20 (not connected, accepting connect from any host)
id=21 (not connected, accepting connect from any host)
id=22 (not connected, accepting connect from any host)
id=23 (not connected, accepting connect from any host)
id=24 (not connected, accepting connect from any host)
id=25 (not connected, accepting connect from any host)
id=26 (not connected, accepting connect from any host)
```

This example shows that all the MySQL NDB nodes are running. If any node is not running, try to restart it. Refer to the [Section MySQL NDB Cluster in IPWorks Troubleshooting Guideline](#) for details.

- If any of the MySQL NDB nodes (Management Node, Data Node, SQL Node) still cannot be started in one SC, go to Step 2.
- If MySQL NDB nodes in both SC still cannot be started, go to Step 3.

2. Recover MySQL NDB in one SC.

- a. Log on to SC where the MySQL NDB nodes cannot start. Meanwhile MySQL NDB nodes in another SC are running.
- b. Stop Storage Server in the SC.

```
# ipw-ctr stop ss
```

If both Storage Server processes are started in both SCs with Active-Standby mode and Storage Server is “Active” in this SC, stopping the Storage Server will activate it in another SC.

- c. Stop all MySQL NDB nodes in this SC that have the problem.

```
# /etc/init.d/ipworks.mysql stop
```

- d. Recover MySQL NDB in this SC.

```
# /etc/init.d/ipworks.mysql recover
```

After this step, the MySQL NDB in the problematic SC shall synchronize with MySQL NDB in another healthy SC and recover the status.



If the MySQL NDB can start in the previous problematic SC. You can continue to perform IPWorks User Data restore that includes NDB data. Refer to [Restore Backup](#) for details.

If the problem is not resolved, perform Step 3 to recover the whole NDB cluster in both SCs.

3. Recover MySQL NDB in both SCs.

Refer to the [Section MySQL NDB Cluster Cannot Work Normally in IPWorks Troubleshooting Guideline](#) to recover whole MySQL NDB cluster to initialization status. You can continue to perform IPWorks User Data restore that includes NDB data running in both SCs.

4.6 IPWorks Application Component Cannot Start

In case certain application is in failure status, use command **ipw-ctr repaired [comp] [<hostname>]** to recover it. If the problem cannot be resolved, go to [Section 5.2 Consult Next Level of Support](#) on page 59. Execute the following recover procedures if next support asks:

1. Reboot the node to recover it, go to [Section 4.6.1 Rebooting PL Node](#) on page 30.
2. Perform power cycle to recover the SC node when SC OS is available, go to [Section 4.6.2 SC Node Power Cycle Graceful](#) on page 31.
3. Perform health check. Refer to [IPWorks Manual Health Check](#) for details.

4.6.1 Rebooting PL Node

For the unknown reason that the IPWorks application cannot start up, rebooting node would help.

To reboot the SC or PL node, follow the steps below:

1. Reboot the PL node:

```
# lde-reboot -n <node id>
```

For example:

```
#lde-reboot -n 3
```

This command will also reboot PL-3.

2. When the procedure is completed, do the following:
 - In case the problem is resolved, go to [Section 5.1 Problem Solved](#) on page 59.
 - In case the problem is not resolved, go to [Section 5.2 Consult Next Level of Support](#) on page 59.



4.6.2 SC Node Power Cycle Graceful

Usually, you shall use `lde-reboot -n <node id>` command instead of power off and then power on (power cycle) the SC node after IPWorks installation. If you really want to perform node power cycle when the node OS is still available, follow the below steps.

Note: If the node OS is unavailable, only perform Step 4.

1. Log on to the SC that will be powered off.

```
#ssh root@<SC IP>
```

Enter root password.

2. Stop Storage Server.

```
# ipw-ctr stop ss
```

3. Stop MySQL NDB nodes (Management Node, Data Node, and SQL Node).

```
# /etc/init.d/ipworks.mysql stop
```

4. Perform power cycle for node.

Perform Section 4.9 Hard Reboot Instance on page 33 in Atlas GUI or CLI.

5. Check MySQL NDB status after SC node startup.

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

If there is any issue in MySQL NDB, refer to the Section MySQL NDB Cluster in IPWorks Troubleshooting Guideline.

6. Start Storage Server.

```
# ipw-ctr start ss
```

For PL node, this graceful power cycle is unnecessary.

4.7 IPWorks Application Component Patch Fails

Prerequisites:

Before the IPWorks patching, periodical System Data backup and User Data backup (with or without NDB) have been performed. And several backup files are created before the patching.

Perform the following recovery procedures:

1. Restore System Data backup.

User shall not cancel the system restore. Refer to [Restore Backup](#) for details.



Note: System Data restore only reinstalls OS and IPWorks application components, it does not remove MySQL NDB data.

2. Perform IPWorks User Data restore that includes MySQL NDB data. Refer to [Restore Backup](#) for details.

Note:

- You shall not perform any cancel operation during this system restore.
- You shall not start or stop any services during the system restore.
- You shall not reboot or power cycle any node during the system restore.

You should wait for the system restore finished before performing any further actions.

3. Manually start MySQL NDB cluster in both SCs.

In case IPWorks User Data backup file includes NDB data, make sure that MySQL NDB cluster is running in both SCs before restore. Otherwise, the restore fails. Refer to the Section MySQL NDB status in [IPWorks Troubleshooting Guideline](#) to start MySQL NDB cluster. If restore failed, go to Section 4.5.1 Recover MySQL NDB for IPWorks User Data Restore on page 28.

4. In case IPWorks User Data backup file does not include MySQL NDB data, the restore does not impact the existed NDB data. You shall not cancel the User Data restore
5. Perform health check. Refer to [IPWorks Manual Health Check](#) for details.

4.8 Recovering Two Damaged SC VMs

Pre-condition:

Backups which include system backup and userdata backup with ndb, are exported to external storage system.

To recover two damaged SCs:

1. Re-deploy IPWorks, refer to [IPWorks Auto Deployment Guideline for KVM - DL380 Gen9](#), [IPWorks Auto Deployment Guideline for KVM - DL380 Gen10](#), or refer to [IPWorks Deployment Guide \(for CEE\)](#).
2. Import backups from external storage system, refer to [Import Backup](#).
3. If EPC AAA PKI Authentication is used, upload all certificate files to the designated directories on SC-1 according to section [Uploading Certificate Files in Configure EPC AAA](#).



4. Restore system data backup, refer to section Restore System Data Backup in [Restore Backup](#).
5. Restore user data backup, refer to section Restore User Data Backup in [Restore Backup](#).

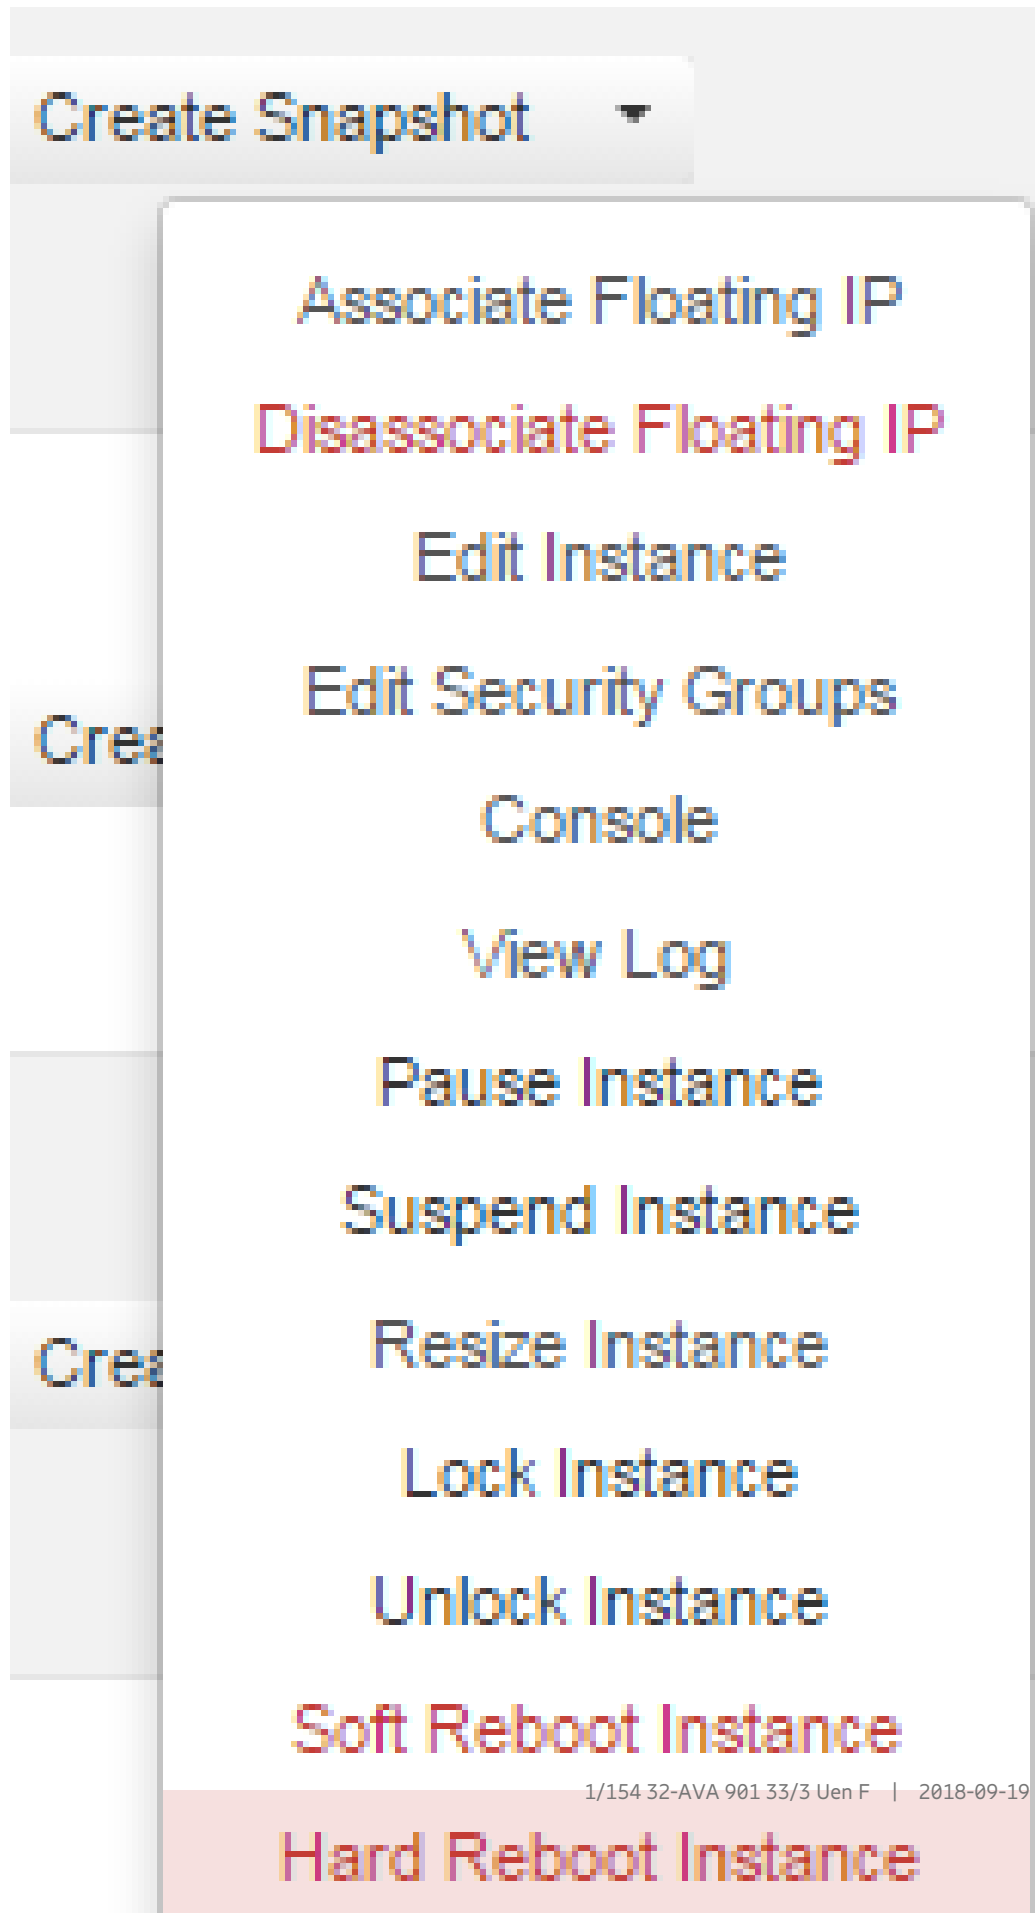
4.9 Hard Reboot Instance

Use one of the methods to hard reboot instance (node):

- Section 4.9.1 By Using Atlas GUI on page 33
- Section 4.9.2 By Using Atlas CLI on page 35

4.9.1 By Using Atlas GUI

User logs on to Atlas GUI and then go to **Project** view, select **instance** in left **Compute** panel, and then in the right panel, search the instance name and then click the **Actions** drop-down list, click **Hard Reboot Instance**. A hard reboot power cycles the instance.





4.9.2 By Using Atlas CLI

User can also execute the operation in CLI in Atlas CLI or CIC CLI.

```
$nova reboot --reboot <instance-id or instance-name>
```

The instance-id or instance-name can be found by using command:

```
$nova list
```

Example:

```
$nova reboot -reboot ipw6a_SC-1
```

```
$nova reboot -reboot 2212c933-71aa-4b92-b53f-3e0081946203
```

4.10 Blade Replacement

This section describes how to replace and recover a GEP blade (SC or PL node blade) when the blade is damaged.

4.10.1 Duration for Blade Replacement

Table 2 records the estimated time during the blade replacement.

Table 2 Estimated Time for Blade Replacement

Replacement Area	Estimated Time (min)	Replacement Period
CEE	60~80	Server replacement
IPWorks (SC node)	5~10 ⁽¹⁾	Storage Server recovery
	5~10	DRBD synchronization
	10 ⁽²⁾	MySQL NDB recovery
IPWorks (PL node)	5~10 ⁽³⁾	PL recovery

(1) This duration is from “heat stack-update is executed” to “Storage Server is running”.

(2) This operation costs several minutes, and actual duration is based on the data size in MySQL NDB database.

(3) The duration is for the general operation not including heavy loaded data happened when the service is starting.

4.10.2 Replacing GEP Blade

To replace the GEP blade physically and recover the system successfully, do:

Replace the server

1. Replace the server in the CEE.



For details, refer to Server Replacement.

Note: To avoid node recovery failure, do the following operation before executing `expandcee` command:

- a Navigate to `tmp` directory on fuel.

```
cd /tmp/
```

- b Download the network file.

```
fuel --env 1 network --download
```

- c Modify the downloaded file (`network_1.yaml`).

Remove NIC part for the compute. For example:

```
compute-0-11:
  if1: eth0
  if2: eth1
  if3: eth6
  if4: eth2
  if5: eth7
  if6: eth5
```

- d Upload the modified file.

```
fuel --env 1 network --upload
```

Inspect the lost node

1. Log in to a healthy SC node, then execute command:

```
SC-<x>:~ # tipc-config -n
```

2. Confirm the lost node type (SC or PL) from the command output. For example:

```
SC-1:~ # tipc-config -n
Neighbors:
<1.1.2>: up
<1.1.3>: down
<1.1.4>: up
<1.1.5>: up
<1.1.6>: up
```

In the example, the IPWorks system includes two SC nodes and four PL nodes, and the corresponding relationship between the label number and the node number, shows as below:

```
<1.1.1>: SC-1
<1.1.2>: SC-2
<1.1.3>: PL-3
```




```
<1.1.4>: PL-4
<1.1.5>: PL-5
<1.1.6>: PL-6
```

The command output shows the status of PL-3 is down, and indicate that the lost node is PL-3.

If no node is shown down status, the lost node is a PL node. Since the PL node evacuates to other compute resource based on ha-offline policy, recover the PL node is not necessary. To check all the compute resources where the PL evacuates to, refer to Step Get the resource name.

Recover the lost node

1. To recover the PL node, refer to Page 37.
2. To recover the SC node, refer to Page 38.

Recover the PL node

Use `nova rebuild` to recover the PL. During the recovery, the PL is boot up from a healthy SC blade, and installs all rpm packages automatically.

1. Connect to Atlas.

```
#ssh atlasadm@<ATLAS_VM_IP_ADDRESS>
atlasadm@atlas:~$ source openrc
```

2. Get the relative information of the lost PL node.

- a. Identify the stack name of the IPWorks.

```
atlasadm@atlas:~$ openstack stack list
```

- b. Get the resource name.

```
atlasadm@atlas:~$ openstack stack resource list
<stack_name>
```

- c. Get the PL image ID.

```
atlasadm@atlas:~$ openstack stack resource show
<stack_name> <resource_name> |grep image
```

In the following example, the image ID is 3637a421-25ea-41f3-b878-8e716a9fff95.

```
atlasadm@atlas:~$ openstack stack resource show sub811_lsv22_0827
ipw_PL-3 |grep image
|                                     |  "image": {
|                                     |      |
|                                     |      "href": "https://⇒
```



```
ipworks.sub8.ctrl.ericsson.se:8774/images/⇒
3637a421-25ea-41f3-b878-8e716a9fff95",
```

3. Rebuild a new PL.

```
atlasadm@atlas:~$ nova rebuild <server> <image>
```

Note: server is the ID number, and can be got from command:
`nova list | grep <stack_name> | grep <Internal IP>`.
 Internal IP is the internal IP of the lost PL node, and can be got in
 /etc/hosts on a healthy SC with the PL node confirmed in Step 2.

image is the image ID got from Step c.

The following is an example of how to get the server ID number.

```
atlasadm@atlas:~$
nova list | grep sub811_lsv22_0827 | grep 169.254.100.3
| 08516d2f-9dae-4c33-9f92-52767c2557ea | sub811_lsv22_0827_⇒
PL-3 | ACTIVE | - | Running | ⇒
sub811_lsv22_0827_int_sp=169.254.100.3; ⇒
sub811_lsv22_0827_sig_sp=192.168.15.3; ⇒
sub811_lsv22_0827_data_sp=192.168.16.3 |
```

The following is an example of how to rebuild a new PL.

```
atlasadm@atlas:~$
```

```
nova rebuild 08516d2f-9dae-4c33-9f92-52767c2557ea
3637a421-25ea-41f3-b878-8e716a9fff95
```

4. Wait for several minutes, then check the status of the recovered PL.

- a. Log in to a healthy SC.
- b. To check the status of the recovered PL is up, execute commands:

```
SC-<x>:~ # tipc-config -n
```

- c. To check the service is running on the recovered PL, execute commands:

```
SC-<x>:~ # ipw-ctr status all
```

Recover the SC node

Use `heat stack-update` to recover the SC with the new SC image and yaml file.

Note: The environment file (`env_file.yaml`) is saved after the IPWorks deployment process.

1. Connect to Atlas.



```
#ssh atlasadm@<ATLAS_VM_IP_ADDRESS>
```

```
atlasadm@atlas:~$ source openrc
```

2. Get image `sc-pxeboot.qcow2`.

- a. Get IPW1.9 (or latter) package from GW link.
- b. Upload the package to directory: `/tmp`.
- c. Uncompress the package (shown as below), and find image `sc-pxeboot.qcow2` in `/images` directory.

```
atlasadm@atlas:~$ cd /tmp
atlasadm@atlas:~$ tar -xzf <package>
atlasadm@atlas:~$ cd images/
atlasadm@atlas:~$ ls
```

3. Create SC image based on `sc-pxeboot.qcow2`.

```
atlasadm@atlas:~$
glance image-create --name <sc_pxeboot_image_name>
--disk-format qcow2 --container-format bare --file
sc-pxeboot.qcow2
```

Note: `sc_pxeboot_image_name` is a specific name used for GEP replacement process.

The following is an example of how to create a new SC image:

```
atlasadm@atlas:~$
glance image-create --name SC_pxeboot --disk-format qcow2
--container-format bare --file /tmp/images/sc-pxeboot.qcow2
```

4. Modify the HOT yaml file.

- a. Get the image ID of the image name (`sc_pxeboot_image_name`).

```
atlasadm@atlas:~$ glance image-list | grep
<sc_pxeboot_image_name>
```

In the following example, the image ID is `55e4b5d7-3c78-4fce-b4b0-70fc9d4bff94`.

```
atlasadm@atlas:~$ glance image-list |grep SC_pxeboot
| 55e4b5d7-3c78-4fce-b4b0-70fc9d4bff94 | SC_pxeboot
```

- b. Get the active HOT yaml file.

```
atlasadm@atlas:~$ openstack stack template show
<stack_id> > ipw_hot_onboarding.yaml
```



- c. Modify the active HOT yaml file.

In the yaml file, replace the image param of the damaged resource (for example, ipw_SC-1) with the image ID got from Step a. In the following example, the lost node is SC-2:

```
ipw_SC-2:
  ...
  properties:
    config_drive: 'True'
    flavor:
      get_param: SC_FLAVOR_NAME
    image:
      get_param: SC_IMAGE
  ...
```

Replace the image param with the image ID.

```
ipw_SC-2:
  ...
  properties:
    config_drive: 'True'
    flavor:
      get_param: SC_FLAVOR_NAME
    image: 55e4b5d7-3c78-4fce-b4b0-70fc9d4bff94
  ...
```

5. Update the stack to rebuild the new SC.

```
atlasadm@atlas:~$ heat stack-update -f ipw_hot_onboarding.yaml
-e <env_file.yaml> <stack_name> --rollback true
```

6. Keep checking the status until UPDATE_COMPLETE is shown.

```
atlasadm@atlas:~$ openstack stack event list <stack_name>
```

UPDATE_COMPLETE means the recovery is successfully.

7. Log in to an SC node and check the DRBD synchronization status.

```
SC-<x>:~ # drbd-overview
```

The DRBD can synchronize automatically. In the following example, the synchronization process is 3.5%.

```
SC-1:~ # drbd-overview
0:drbd0/0 SyncSource Primary/Secondary UpToDate/⇒
Inconsistent C r----- lvm-pv: lde-cluster-vg 100.00g 50.06g
[>.....] ⇒
sync'ed: 3.5% (98916/102400)M
```



The following shows that the synchronization is completed.

```
SC-x:~ # drbd-overview
0:drbd0/0 Connected Primary/Secondary UpToDate/UpToDate =>
C r----- lvm-pv: lde-cluster-vg 100.00g 50.06g
```

8. Resize disk and reboot on SC node.

```
# ssh root@<MIP_OAM_IP>
```

```
SC-X:~# ipworks.mysql stop-ndbcluster
```

```
SC-X:~# umount /local/ipworks
```

```
SC-X:~# /opt/ipworks/common/scripts/ipwResizePartition.sh -d
/dev/vda -p 7
```

Note: If you receive the following error message, this means that the new partition cannot be used immediately, this issue will be resolved after SC reboot.

```
Calling ioctl() to re-read partition table.
Re-reading the partition table failed.: Device or resource busy
```

```
SC-X:~# reboot
```

Resize partition on both SC-X after reboot is done.

```
# ssh root@<MIP_OAM_IP>
```

```
SC-X:~# resize2fs /dev/vda7
```

```
SC-X:~# ipworks.mysql start-ndbcluster
```

9. Recover the MySQL NDB in the recovered SC blade.

- a. Log in to the recovered SC.
- b. Stop the Storage Server and MySQL.

```
SC-X:~# ipw-ctr stop ss
```

```
SC-X:~# /etc/init.d/ipworks.mysql stop
```

- c. Recover MySQL NDB.

```
SC-X:~# /etc/init.d/ipworks.mysql recover
```

After the recovery, the MySQL can be logged in, and all the data in /local/ipworks directory is restored.



- d. Check the status of all the nodes on SCs.

```
SC-X:~# /etc/init.d/ipworks.mysql show-status
```

For example:

```
SC-2:~ # /etc/init.d/ipworks.mysql show-status
Connected to Management Server at: localhost:1186
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=27   @169.254.100.1  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0, *)
id=28   @169.254.100.2  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0)
[ndb_mgmd(MGM)] 2 node(s)
id=1     @169.254.100.1  (mysql-5.6.31 ndb-7.4.12)
id=2     @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
[mysqld(API)]   24 node(s)
id=3 (not connected, accepting connect from SC-1)
id=4     @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
id=5 (not connected, accepting connect from any host)
id=6 (not connected, accepting connect from any host)
id=7 (not connected, accepting connect from any host)
id=8 (not connected, accepting connect from any host)
```

Note: The expected result is that at least the nodes with ID value 1, 2, 27 and 28, are started, and one of the nodes with ID value 3 or 4, is started.

- e. Restart Storage Server on the new SC.

```
SC-X:~# ipw-ctr restart ss
```

4.10.3 Heath Check

Perform health check activities, refer to [IPWorks Auto Health Check](#) or [IPWorks Manual Health Check](#) for details.

4.10.4 Create a Backup

Create a backup after the replacement, refer to [Create Backup](#).

4.11 IPWorks environment file, images and yaml file missing

Perform the following recovery procedures:

1. Log on Atlas server.

```
#ssh atlasadm@<Atlas_addr>
```




```

parameters:
  VNF_NAME: ipw-Crux-1-6
  SC_IMAGE: ipw-Crux-1-6-sc
  PL_IMAGE: ipw-Crux-1-6-pl
  SC_FLAVOR_NAME: ipw-Crux-1-6-sc_0314
  PL_FLAVOR_NAME: ipw-Crux-1-6-pl_0314
  TIME_ZONE: Asia/Shanghai
  SC1_HYPERVISOR_HOSTNAME: nova:compute
  SC2_HYPERVISOR_HOSTNAME: nova:compute

```

5. Upload images and yaml file from IPWorks VNF package to /home/atlasadm/ipworks, then unpack. For detailed information, refer to section Transfer IPWorks VNF Package to Atlas Server in IPWorks Deployment Guide.

4.12 Data Node Can Not Start Up During CEE Upgrade

Perform the following to recovery procedures:

1. Find the ndb related processes on both SCs.

For example:

```
SC-X:~ # ps -ef |grep ndb |grep -v grep
```

```

root 22728 1 0 Aug28 ? 00:10:19 /opt/ipworks/mysql/mysql/sbin/ndb_mgmd -f /home/ipworks/mysql/
root 22977 1 0 Aug28 ? 00:00:06 /opt/ipworks/mysql/mysql/sbin/ndbmtd --defaults-file=/home/ipw
root 22979 22977 3 Aug28 ? 00:38:19 /opt/ipworks/mysql/mysql/sbin/ndbmtd --defaults-file=/home

```

2. Kill all the processes shown in Step 1 on both SCs.

For example:

```
SC-X:~ # kill -9 22728
```

```
SC-X:~ # kill -9 22977
```

```
SC-X:~ # kill -9 22979
```

3. Start Mysql NDB Management node on both SCs.

```
SC-X:~ # /etc/init.d/ipworks.mysql start-mgmd
```

4. Start Mysql NDB Cluster

```
SC-X:~ # /etc/init.d/ipworks.mysql start-ndbcluster
```




5. Check the Mysql NDB Cluster status.

```
SC-X:~ # /etc/init.d/ipworks.mysql show-status
```

4.13 Recovering One Damaged SC VM

If one out of two SCs is damaged, recover the SC VM with command line first, refer to Section 4.13.1 on page 45. If failed, recover it with deployment script, refer to Section 4.13.2 on page 50.

4.13.1 Recovering SC VM with Command Line

1. Connect to the host where the damaged SC VM locates (take SC-2 as an example).

```
#ssh root@<HOST_IP_ADDRESS>
```

2. Get image (sc-pxeboot.qcow2) from IPWorks release package.
 - a. Upload the package to a directory (for example, /tmp) on the host.

The package name is: 19010-CXP9023809_3_Ux_<Revision Number>.tar.gz.

- b. Decompress the package, show as below:

```
#cd /tmp
#tar -zxvf 19010-CXP9023809_3_Ux_<Revision Number>.tar.gz
#cd images
#ls -la
-rw-r--r-- 1 292374 16342 3720151040 Mar 6 19:22 ipw-sc-22.qcow2
-rwxr-xr-x 1 292374 16342 786432 Mar 6 19:22 pxeboot.qcow2
-rwxr-xr-x 1 292374 16342 786432 Mar 6
19:22 sc-pxeboot.qcow2
```

3. Check the instance name of the damaged SC.

```
# virsh list --all
Id      Name                                State
-----
1       IPW-PL-4                            running
2       IPW-SC-2                            running
```

4. Record the storage folder and name of the damaged SC image.

In the following output example, /root/auto_deployment/images/IPW/run/ stands for the storage folder, and IPW-ipw-sc-22.qcow2 stands for the image name.



```
#virsh edit IPW-SC-2
<devices>
  <emulator>/usr/bin/qemu-kvm</emulator>
  <disk type='file' device='disk'>
    <driver name='qemu' type='qcow2' />
    <source file='/root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2' />
    <target dev='vda' bus='virtio' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0' />
  </disk>
</devices>
```

5. Destroy the damaged image.

```
# virsh destroy IPW-SC-2
```

The instance name (IPW-SC-2) is got from Step 3.

6. Navigate to the folder recorded in Step 4.

```
# cd /root/auto_deployment/images/IPW/run/
```

7. Delete or move the damaged image recorded in Step 4.

If delete:

```
# rm -rf /root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2
```

If move:

```
# mv /root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2 /tmp/
```

8. Copy the new image (in Step 2) to the folder recorded in Step 4, and rename it to the name recorded in Step 4.

```
# cp /tmp/images/sc-pxeboot.qcow2 /root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2
```

9. Re-size the qcow2 file virtual size if needed.

- a. Log on to a healthy SC to check /dev/vda size.

```
#ssh root@<healthy SC IP Address>
#fdisk -l
SC-1:~ # fdisk -l
Disk /dev/vda: 280 GiB, 300647710720 bytes, 587202560 sectors
Units: sectors of 1 * 512 = 512 bytes
```

- b. Re-size the qcow2 virtual size to 280 G on host.



```
#cd /root/auto_deployment/images/IPW/run/
#qemu-img info /root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2
image: IPW-ipw-sc1-22.qcow2
file format: qcow2
virtual size: 80G (300647710720 bytes)
disk size: 137G
cluster_size: 65536
Format specific information:
    compat: 1.1
    lazy refcounts: false
    refcount bits: 16
    corrupt: false
```

If the virtual size in the preceding output is not equal to the output by `fdisk` command in Step a, execute the following `fdisk` command to re-size the qcow2 file virtual size. Otherwise, go to Step 10.

Example output by `fdisk` command with 280 G virtual size:

```
#qemu-img resize /root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2 280G
#qemu-img info /root/auto_deployment/images/IPW/run/IPW-ipw-sc2-22.qcow2
image: IPW-ipw-sc1-22.qcow2
file format: qcow2
virtual size: 280G (300647710720 bytes)
disk size: 137G
cluster_size: 65536
Format specific information:
    compat: 1.1
    lazy refcounts: false
    refcount bits: 16
    corrupt: false
```

10. Start the damaged SC on host.

```
#virsh start IPW-SC-2
```

The instance name (IPW-SC-2) is got from Step 3.

11. Log on to an SC node and check the DRBD synchronization status.

```
SC-<x>:~ # drbd-overview
```

The DRBD is synchronized automatically.

Example output 1:

```
SC-1:~ # drbd-overview
0:drbd0/0 SyncSource Primary/Secondary UpToDate/Inconsistent C r----- lvm-pv: lde-clu
```



When the status becomes UpToDate/UpToDate in example output 2, and the second part values in example output 3 are all 0, the synchronization is completed.

Example output 2:

```
SC-x:~ # drbd-overview
0:drbd0/0 Connected Primary/Secondary UpToDate/UpToDate C r----- lvm-pv: lde-cluster-vg 100
```

Example output 3:

```
SC-x:~ # drbdadm get-gi drbd0
BF4BD24D78025EAF:0000000000000000:BD26BDDF3391E066:BD25BDDF3391E066:1:1:1:1:0:0:0
```

12. Log on to the recovered SC and re-size ipworks disk(/dev/vda7) size.

Skip this step if Step 9 is not executed.

The following example, shows /dev/vda7 size before re-size.

```
#ssh root@<recover SC IP Address>
#fdisk -l
SC-1:~ # fdisk -l
Disk /dev/vda: 280 GiB, 300647710720 bytes, 587202560 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: 690E70A6-8A28-49AC-AF10-927E7FA36FCE
```

Device	Start	End	Sectors	Size	Type
/dev/vda1	2048	8390655	8388608	4G	EFI System
/dev/vda2	8390656	29362175	20971520	10G	Microsoft basic data
/dev/vda3	29362176	37750783	8388608	4G	Microsoft basic data
/dev/vda4	37750784	48236543	10485760	5G	Microsoft basic data
/dev/vda5	48236544	121636863	73400320	35G	Microsoft basic data
/dev/vda6	121636864	121899007	262144	128M	Microsoft basic data
/dev/vda7	121899008	142870527	20971520	10G	Microsoft basic data
/dev/vda8	1024	2047	1024	512K	BIOS boot

The following example, shows /dev/vda7 size becomes 221.9 G. The amount of all vda partitions is 280 G, which means re-sizing is successful.



```
#umount /local/ipworks
#/opt/ipworks/common/scripts/ipwResizePartition.sh -d /dev/vda -p 7
#reboot
#ssh root@<recover SC IP Address>
#resize2fs /dev/vda7
#fdisk -l
SC-1:~ # fdisk -l
Disk /dev/vda: 280 GiB, 300647710720 bytes, 587202560 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: 690E70A6-8A28-49AC-AF10-927E7FA36FCE
```

Device	Start	End	Sectors	Size	Type
/dev/vda1	2048	8390655	8388608	4G	EFI System
/dev/vda2	8390656	29362175	20971520	10G	Microsoft basic data
/dev/vda3	29362176	37750783	8388608	4G	Microsoft basic data
/dev/vda4	37750784	48236543	10485760	5G	Microsoft basic data
/dev/vda5	48236544	121636863	73400320	35G	Microsoft basic data
/dev/vda6	121636864	121899007	262144	128M	Microsoft basic data
/dev/vda7	121899008	587202526	465303519	221.9G	Linux filesystem
/dev/vda8	1024	2047	1024	512K	BIOS boot

13. Recover the MySQL NDB in the recovered SC node.

- a. Log on to the recovered SC.

```
#ssh root@<recover SC IP Address>
```

- b. Stop the Storage Server and MySQL.

```
SC-<x>:~ # ipw-ctr stop ss
```

```
SC-<x>:~ # /etc/init.d/ipworks.mysql stop
```

- c. Recover MySQL NDB.

```
SC-<x>:~ # /etc/init.d/ipworks.mysql recover
```

Note: After the recovery, the MySQL can be logged in, and all the data in /local/ipworks directory is restored.

- d. Check the status of all the nodes on SCs.

```
SC-<x>:~ # /etc/init.d/ipworks.mysql show-status
```

For example:



```

SC-2:~ # /etc/init.d/ipworks.mysql show-status
Connected to Management Server at: localhost:1186
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=27   @169.254.100.1  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0, *)
id=28   @169.254.100.2  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0)
[ndb_mgmd(MGM)] 2 node(s)
id=1     @169.254.100.1  (mysql-5.6.31 ndb-7.4.12)
id=2     @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
[mysqld(API)]   24 node(s)
id=3 (not connected, accepting connect from SC-1)
id=4     @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
id=5 (not connected, accepting connect from any host)
id=6 (not connected, accepting connect from any host)
id=7 (not connected, accepting connect from any host)
id=8 (not connected, accepting connect from any host)

```

Note: The expected result is that at least the nodes with ID value 1, 2, 27 and 28, are started, and one of the nodes with ID value 3 or 4, is started.

- e. Restart Storage Server on the new SC.

```
SC-<x>:~ # ipw-ctr restart ss
```

- f. Perform health check, refer to [IPWorks Auto Health Check](#).

4.13.2 Recovering SC VM with Deployment Script

1. Connect to HOST1 where the script and image are deployed.

```
#ssh root@<HOST_IP_ADDRESS>
```

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

```
#cp sc-pxeboot.qcow2 ./<VNF_NAME>
```

```
#!/root/auto_deployment/IPW2/kvm_deployment
```

Note: Make sure file `./config/ipwenv.conf` exists.

2. Clean up the damaged SC VM.

Example for SC-1 cleanup in HOST1:

```
#!/ipwdeploy.sh -a cleanup -l 1 -m 1 -T d
```




```
SC-x:~ # drbdadm get-gi drbd0
0:drbd0/0 Connected(2*) Second/Primar UpToDa/UpToDa
```

Example output 2:

```
SC-x:~ # drbdadm get-gi drbd0
309FE992B4DBABB7:0000000000000000:3F6F408DB1A4AC24:5A6
3078BAB19F526:1:1:1:0:0:0:1:0:0:0
```

5. Re-size ipworks disk (/dev/vda7) size.

```
./ipwInit.sh -c config/ipwenv.conf -l 1
```

```
.....
The filesystem is already 58162939 blocks long. Nothing to do!
20180321-06:13:32: *****resize partition 10.170.15.130 done*****
20180321-06:13:32: IPWorks Init done...
```

6. Recover the MySQL NDB in the recovered SC node.

- a. Log on to the recovered SC.

```
#ssh root@<recover SC IP Address>
```

- b. Stop the Storage Server and MySQL.

```
SC-<x>:~ # ipw-ctr stop ss
```

```
SC-<x>:~ # /etc/init.d/ipworks.mysql stop
```

- c. Recover MySQL NDB.

```
SC-<x>:~ # /etc/init.d/ipworks.mysql recover
```

Note: After the recovery, the MySQL can be logged in, and all the data in /local/ipworks directory is restored.

- d. Check the status of all the nodes on SCs.

```
SC-<x>:~ # /etc/init.d/ipworks.mysql show-status
```

For example:



```
SC-2:~ # /etc/init.d/ipworks.mysql show-status
Connected to Management Server at: localhost:1186
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=27   @169.254.100.1  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0, *)
id=28   @169.254.100.2  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0)
[ndb_mgmd(MGM)] 2 node(s)
id=1    @169.254.100.1  (mysql-5.6.31 ndb-7.4.12)
id=2    @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
[mysqld(API)]   24 node(s)
id=3 (not connected, accepting connect from SC-1)
id=4    @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
id=5 (not connected, accepting connect from any host)
id=6 (not connected, accepting connect from any host)
id=7 (not connected, accepting connect from any host)
id=8 (not connected, accepting connect from any host)
```

Note: The expected result is that at least the nodes with ID value 1, 2, 27 and 28, are started, and one of the nodes with ID value 3 or 4, is started.

- e. Restart Storage Server on the new SC.

```
SC-<x>:~ # ipw-ctr restart ss
```

- f. Perform health check, refer to [IPWorks Auto Health Check](#).

4.14 Recovering One Damaged PL VM

If one out of two PLs is damaged, to recover the PL VM:

1. Connect to HOST1 where the script and image are deployed.

```
# ssh root@<HOST_IP_ADDRESS>
# cd /root/auto_deployment/images
# cp ./pxeboot.qcow2 ./<VNF_NAME>
# cd /root/auto_deployment/IPW2/kvm_deployment
```

Note: Make sure file `./config/ipwenv.conf` exists.

2. Clean up the damaged PL VM.

Example for PL-3 cleanup in HOST1:

```
#./ipwdeploy.sh -a cleanup -l 1 -m 3 -T d
```



```
DHOST1 [*. *.*]: Try Delete VM [ IPW-PL-3 ] OVS ports and stop VM.
DHOST1 [*. *.*]: Start to clean related libvirt xml.
```

3. Deploy a new PL VM.

Example for PL-3 deployment in HOST1:

```
#./ipwdeploy.sh -a deploy -l 1 -m 3 -T d
```

```
DHOST1 [*. *.*]: Precheck Environment: OS configuration
DHOST1 [*. *.*]: Precheck Environment: Existed Virtual Machine(s)
DHOST1 [*. *.*]: Copy PL template image to running folder
DHOST1 [*. *.*]: Copy PL template image to running folder
DHOST1 [*. *.*]: Generate libvirt xml files
DHOST1 [*. *.*]: Prepare Environment: OVS and DPDK
DHOST1 [*. *.*]: Prepare Environment: Qemu
DHOST1 [*. *.*]: Start Virtual Machine: IPW-PL-3
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
[DHOST1 : *.*.*]
*****
VM status
*****
  Id      Name              State
-----
  11      IPW-PL-3              running

*****
OVS Brige and Port Status
*****
Bridge [ br-int ]
      Port "IPW-PL-3-eth0"
Bridge [ br-trf ]
      Port "IPW-PL-3-eth1"
      Port "IPW-PL-3-eth2"

*****
Libvirt Status
*****
Active: active (running) since Thu 2018-05-17 13:38:00 CST; 3h 56min ago

*****
ovs-network Status
*****
  Name              State    Autostart    Persistent
-----
  IPW-ovs-network    active   yes          yes
```

4.15 Recovering SC VM and PL VM in One Host

If one host is damaged, to replace the SC VM and PL VM:

1. Connect to HOST1 where the script and image are deployed.

```
# ssh root@<HOST_IP_ADDRESS>
# cd /root/auto_deployment/images
# cp sc-pxeboot.qcow2 ./<VNF_NAME>
```



```
# cp ./pxeboot.qcow2 ./<VNF_NAME>
```

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

Note: Make sure file ./config/ipwenv.conf exists.

2. Clean up the damaged SC VM and PL VM.

Example for SC-1 and PL-3 cleanup in HOST1:

```
#./ipwdeploy.sh -a cleanup -l 1 -T d
```

```
DHOST1 [*.*.*]: Try Delete VM [ IPW-SC-1 ] OVS ports and stop VM.
DHOST1 [*.*.*]: Try Delete VM [ IPW-PL-3 ] OVS ports and stop VM.
DHOST1 [*.*.*]: Start to clean related ovs-network
DHOST1 [*.*.*]: Start to clean related ovs ifaces and bridges
DHOST1 [*.*.*]: Start to clean related libvirt xml.
DHOST1 [*.*.*]: Start to clean SC-1 qcow2 files if needed.
```

3. Deploy new SC VM and PL VM in one host.

Example for SC-1 and PL-3 deployment in HOST1:

```
#./ipwdeploy.sh -a deploy -l 1 -T d -r
```



```

DHOST1 [*. *.*]: Precheck Environment: OS configuration
DHOST1 [*. *.*]: Precheck Environment: Existed Virtual Machine(s)
DHOST1 [*. *.*]: Copy SC template image to running folder
DHOST1 [*. *.*]: Copy SC template image to running folder
DHOST1 [*. *.*]: Inject parameters to SC image.
Done
DHOST1 [*. *.*]: Generate 1 SC image from injected parameters SC image
DHOST1 [*. *.*]: Generate SC-1 image from injected parameters SC image
DHOST1 [*. *.*]: Generate libvirt xml files
DHOST1 [*. *.*]: Prepare Environment: OVS and DPDK
DHOST1 [*. *.*]: Prepare Environment: Qemu
DHOST1 [*. *.*]: Start Virt Network
DHOST1 [*. *.*]: Start Virtual Machine: IPW-SC-1 (wait 60 sec)
DHOST1 [*. *.*]: Start Virtual Machine: IPW-PL-3
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
[DHOST1 : *.*.*]
*****
VM status
*****
Id      Name      State
-----
12      IPW-SC-1    running
13      IPW-PL-3    running

*****
OVS Brige and Port Status
*****
Bridge [ br-int ]
Port "IPW-PL-3-eth0"
Port "IPW-SC-1-eth0"
Bridge [ br-oam ]
Port "IPW-SC-1-eth2"
Port "IPW-SC-1-eth1"
Port "IPW-PL-3-eth2"
Port "IPW-PL-3-eth1"
Bridge [ br-trf ]
Port "IPW-SC-1-eth2"
Port "IPW-SC-1-eth1"
Port "IPW-PL-3-eth2"
Port "IPW-PL-3-eth1"

*****
Libvirt Status
*****
Active: active (running) since Thu 2018-05-17 13:38:00 CST; 3h 59min ago

*****
ovs-network Status
*****
Name      State      Autostart    Persistent
-----
IPW-ovs-network  active     yes          yes

```

4. Log on to an SC node and check the DRBD synchronization status.

When the status becomes UpToDate/UpToDate in example output 1, and the second UUID block values in example output 2 are all 0, the synchronization is completed.

Example output 1:

```

SC-x:~ # drbd-overview
0:drbd0/0 Connected(2*) Second/Primar UpToDa/UpToDa

```



Example output 2:

```
SC-x:~ # drbdadm get-gi drbd0
309FE992B4DBABB7:0000000000000000:3F6F408DB1A4AC24:5A6
3078BAB19F526:1:1:1:0:0:0:1:0:0:0
```

5. Re-size ipworks disk (/dev/vda7) size.

```
./ipwInit.sh -c config/ipwenv.conf -l 1
```

```
.....
The filesystem is already 58162939 blocks long. Nothing to do!
20180321-06:13:32: *****resize partition 10.170.15.130 done*****
20180321-06:13:32: IPWorks Init done...
```

6. Recover the MySQL NDB in the recovered SC node.

- a. Log on to the recovered SC.
- b. Stop the Storage Server and MySQL.

```
SC-<x>:~ # ipw-ctr stop ss
```

```
SC-<x>:~ # /etc/init.d/ipworks.mysql stop
```

- c. Recover MySQL NDB.

```
SC-<x>:~ # /etc/init.d/ipworks.mysql recover
```

Note: After the recovery, the MySQL can be logged in, and all the data in /local/ipworks directory is restored.

- d. Check the status of all the nodes on SCs.

```
SC-<x>:~ # /etc/init.d/ipworks.mysql show-status
```

For example:



```
SC-2:~ # /etc/init.d/ipworks.mysql show-status
Connected to Management Server at: localhost:1186
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=27   @169.254.100.1  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0, *)
id=28   @169.254.100.2  (mysql-5.6.31 ndb-7.4.12, =>
Nodegroup: 0)
[ndb_mgmd(MGM)] 2 node(s)
id=1     @169.254.100.1  (mysql-5.6.31 ndb-7.4.12)
id=2     @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
[mysqld(API)]   24 node(s)
id=3 (not connected, accepting connect from SC-1)
id=4     @169.254.100.2  (mysql-5.6.31 ndb-7.4.12)
id=5 (not connected, accepting connect from any host)
id=6 (not connected, accepting connect from any host)
id=7 (not connected, accepting connect from any host)
id=8 (not connected, accepting connect from any host)
```

Note: The expected result is that at least the nodes with ID value 1, 2, 27 and 28, are started, and one of the nodes with ID value 3 or 4, is started.

- e. Restart Storage Server on the new SC.

```
SC-<x>:~ # ipw-ctr restart ss
```

- f. Perform health check, refer to IPWorks Auto Health Check.



5 Problem Reporting

In general, all the described recovery situations must be seen as abnormal and must be reported to the next level of support or according to other documented procedure such as log book, even if the recovery has been successful. Often a Customer Service Request (CSR) is written to a responsible support organization.

If the situation has affected the ISP, it must be reported as such according to documented procedure.

In many situations, it is required to perform a Root Cause Analysis (RCA) afterwards to determine the source of the problem. It is therefore important to carefully document the problematic situation and all the recovery steps that have been taken.

Many log files in the system must be saved or copied to another place to prevent them from being overwritten with newer information. It is important that these logs are available for any future RCA.

For information about how to collect data and log files, refer to [Data Collection Guideline for IPWorks](#).

5.1 Problem Solved

The recovery seems to have worked. Keep the site and the affected functions under extra observation for a while to ensure that the fault does not reoccur.

Record the incident according to local procedures using a log book or similar.

5.2 Consult Next Level of Support

Provide the receiving support organization with the following information:

- Site name
- Location
- Operator
- Contact information
- Node version loaded
- IS version or TSP version loaded on the node
- Node configuration
- Problem description



- Procedures followed, including document number
- Information about activities done before, during, and after the incident
- Logs

For information about how to collect data and log files, refer to [Data Collection Guideline for IPWorks](#).



6 Appendix

Interpretation for the site-specified parameters used in commands:

1 Parameters in script file `./ipwdeploy.sh`:

- `-l <hostlist>` stands for host, for example, `-l 1` means DHOST1, and `-l 2` means DHOST2 in `config/ipwenv.conf`.
- `-m <vmlist>` stands for VM, for example, `-m 1` means SC-1, `-m 2` means SC-2, `-m 3` means PL-3, and `-m 4` means PL-4.
- `-T type` stands for type, for example, `-T d` means the type of the Host is Deploy.
- `-r` stands for FLAG for recovering single VM.

For more parameter explanations, check by executing commands:
`./ipwdeploy.sh --help`.

2 Parameters in script file `./ipwInit.sh`:

- `-c conf` stands for main configuration file for IPWorks VNF and BSP configuration (for example, `ipwenv.conf`).
- `-l <hostlist>` stands for specified host, `-l 1` means DHOST1, and `-l 2` means DHOST2 in `config/ipwenv.conf`.

For more parameter explanations, check by executing commands:
`./ipwInit.sh --help`.





Reference List

Ericsson Documents

- [1] Glossary of Terms and Acronyms
- [2] Trademark Information
- [3] Typographic Conventions
- [4] IPWorks Configuration Management
- [5] IPWorks Troubleshooting Guideline
- [6] Data Collection Guideline for IPWorks
- [7] Backup and Restore
- [8] Personal Health and Safety Information
- [9] System Safety Information
- [10] IPWorks Auto Health Check
- [11] IPWorks Manual Health Check
- [12] Configure MySQL NDB Cluster
- [13] Create Backup
- [14] Restore Backup
- [15] Check Alarm Status
- [16] Configure EPC AAA
- [17] IPWorks Deployment Guide, 21/1553-AVA 901 33/3 Uen
- [18] Emergency Recovery Procedure, 2/154 32-AZE 102 01 Uen
- [19] Server Replacement, 4/1543-CSA 113 125/4 Uen