NEC

SX-Aurora TSUBASA

SX-Aurora TSUBASA

Installation Guide

Proprietary Notice

The information disclosed in this document is the property of NEC Corporation (NEC) and/or its licensors. NEC and/or its licensors, as appropriate, reserve all patent, copyright, and other proprietary rights to this document, including all design, manufacturing, reproduction, use and sales rights thereto, except to the extent said rights are expressly granted to others.

The information in this document is subject to change at any time, without notice.

Trademarks and Copyrights

- Linux is a registered trademark of Linus Torvalds in the United States and other countries.
- Red Hat and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the United States and other countries.
- Apache is a registered trademark of Apache Software Foundation.
- InfiniBand is a trademark or service mark of InfiniBand Trade Association.
- NVIDIA is a trademark and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries.
- Windows are registered trademarks of Microsoft Corporation in the United States and other countries.
- All other product, brand, or trade names used in this publication are the trademarks or registered trademarks of their respective trademark owners.

© 2018 NEC Corporation

Preface

This document explains how to install, configure, update, and uninstall the SX-Aurora TSUBASA software on the SX-Aurora TSUBASA system.

The latest version of this document is available at:

https://sxauroratsubasa.sakura.ne.jp/documents/guide/pdfs/InstallationGuide_E.pdf

"SX-Aurora TSUBASA Setup Guide" is also available at the following URL, and explains how to set up the SX-Aurora TSUBASA system for first-time users, including hardware setup, installation of the OS and SX-Aurora TSUBASA software, basic environment settings, and execution of sample programs:

https://sxauroratsubasa.sakura.ne.jp/documents/guide/pdfs/SetupGuide_E.pdf

Please note that the setup guide explains the setup procedures mostly for SX-Aurora TSUBASA Model A100-1, and does not describe installation of ScaTeFS and NQSV.

⚠ Note

Please execute command lines starting with "#" as the superuser in this document.

Definitions and Abbreviations

Term	Description
Vector Engine (VE)	The core part of the SX-Aurora TSUBASA system, on which applications are executed. A VE is implemented as a PCI Express card and attached to a server called a vector host.
Vector Host (VH)	A Linux (x86) server to which VEs are attached, in other words, a host computer equipped with VEs.
Vector Island (VI)	A set of a VH and VEs that are attached to the VH. A VI is the basic unit for the tower model and rack mount model described below.
Tower model	One of the SX-Aurora TSUBASA product models. The tower model is a desk side model that can be simply set-up.
Rack mount model	One of the SX-Aurora TSUBASA product models. The rack mount model is a 1U or 4U server model with a server rack. It covers from small systems to large scale systems.
Supercomputer model	One of the SX-Aurora TSUBASA product models. The supercomputer model is positioned as the next generation model of the SX series. It can mount up to eight 4U rack mount servers. All vector engines have water cooling devices.
VE1	In this document, products with a Vector Engine Type whose tens digit is 1 are represented.
VE2	In this document, products with a Vector Engine Type whose tens digit is 2 are represented.
VE3	In this document, products with a Vector Engine Type whose tens digit is 3 are represented.
VMC	Abbreviation of VE Management Controller
IB	Abbreviation of InfiniBand
HCA	Abbreviation of Host Channel Adapter. A kind of PCIe card to connect a server to an IB network.
MPI	Abbreviation of Message Passing Interface. MPI is a standard specification for a communication library. It can be used together with OpenMP or automatic parallelization.
NEC yum repository	The yum repository for NEC SX-Aurora TSUBASA software. The yum repository for the free software can be accessed by any user. The yum repository for the paid software can be accessed only by users with PP support contract.

Term	Description
License server	A server that manages licenses for the software on the SX-Aurora TSUBASA. This is needed to use the paid software.
Frontend, Frontend machine	In this document, frontend (or frontend machine) means compiling programs for VEs on a machine other than SX-Aurora TSUBASA system (or the compile machine).
PP support	The support services to provide technical support for the purchased software products for a fee.

Contents

List of tables

List of figures

Chapter1 Introduction

1.1 Scope

1.1.1 List of software

This document explains installation, configuration, update, and uninstallation of the SX-Aurora TSUBASA software, which is listed in Table 1.

Table 1 The SX-Aurora TSUBASA Software

Software Name	Description
VEOS	VE management software
MMM	Monitoring & Maintenance Manager
VMC Firmware	VMC Firmware
InfiniBand for SX-Aurora TSUBASA	InfiniBand control software
License Server	License management software
License Access Library	License check library
NEC Software Development Kit for Vector Engine (abbreviation: SDK)	Software Development Software There are 2 kinds - SDK Community Edition (abbreviation: SDK-CE) and SDK Standard Edition (abbreviation: SDK-SE). SDK-CE is free software and is available beginning January 2023 and SDK-SE is paid software and is available prior to January 2023. The model numbers of SDK-SE are the followings. UWAA00-N10E-I UWAA00-N11E-I UWAA00-N12E-I UWAA00-N1UE-I Abbreviation "SDK" refers to both SDK-CE and SDK-SE.
SDK Runtime	The runtime libraries and commands in SDK for executing VE programs and MPI programs.
NEC Compat C++ Standard Library for Vector Engine (abbreviation: Compat C++)	C++ standard library for compatibility with C++ compilers prior to version 3.5.x. This is a library for extension of SDK-CE and paid software. You need to purchase this separately from SDK-CE.
NEC Scalable Technology File System (abbreviation: ScaTeFS)	Scalable Technology File System
NEC Network Queuing System V (abbreviation: NQSV)	Batch Execution System

(*1) ScaTeFS is only available in environments that use InfiniBand for the operation network.

It is not available in environments that use only Ethernet for the operation network.

1.1.2 How to get software

There are 2 types of software for SX-Aurora TSUBASA depending on how to get them. Please refer to Section 1.4.1 for the conditions to access the NEC yum repository with restricted access.

A) Free software. You can install the software packages from the NEC yum repository with free access with the yum command.

B) Paid software. If you have the PP support contract, you can install the software packages from the NEC yum repository with restricted access with the yum command. Otherwise, you can install the packages downloaded in the internet delivery product download service.

Components included in each software and how to get them are listed in Table 2.

Table 2 Software components and How to get

Software Name	Components	How to get
VEOS	VEOS	A
MMM	MMM	A
VMC Firmware	VMC Firmware	A
InfiniBand for SX-Aurora TSUBASA	InfiniBand for SX-Aurora TSUBASA	A
License Server	License server	A
License Access Library	License access library	A
SDK	C/C++ Compiler Fortran Compiler binutils Numeric Library Collection NLCPy mpi4py-ve NEC Parallel Debugger (*) Tuning Tool NEC MPI (including NEC MPI/Scalar-Vector Hybrid (*1))	A (SDK-CE) B (SDK-SE) (*) is only in B
SDK Runtime	The binutils, runtime libraries and MPI execution command in SDK	A
Compat C++	NEC Compat C++ Standard Library for Vector Engine	B
ScaTeFS	ScaTeFS/Client	B
NQSV	NQSV/JobServer NQSV/Client	B

(*1) Installation of NEC MPI/Scalar Vector Hybrid is not necessary because the package for

it is the same package for NEC MPI.

1.2 System Requirement

1.2.1 Hardware

The SX-Aurora TSUBASA is available in the following models.

Table 3 Models

	Tower	Rack Mount				Supercomputer
Model Name(*1)	A100-1 A101-1 A111-1	A300-2	A300-4 A311-4	A300-8 A311-8 B300-8 B302-8	A412-8 B401-8 B403-8 C401-8	A500-64 A511-64
Max. # of Vector Engines (VEs)	1	2	4	8	8	64
# of Vector Hosts (VHs)	1	1	1	1	1	8

(*1) The VE installed in Model Name Axxx-x is Type 1*(VE1). The VE installed in Model Name Bxxx-x is Type 2*(VE2). The VE installed in Model Name Cxxx-x is Type 3*(VE3). Please select the installation package that corresponds to the installed VE.

Please refer to the SX-Aurora TSUBASA catalogue for details.

⚠ Note: Boot mode setting of VHs

The boot mode setting in the BIOS of VHs should be left the UEFI mode, which is the factory default setting. The SX-Aurora TSUBASA does not support other modes.

1.2.2 Supported Operating Systems

The SX-Aurora TSUBASA software runs on the Linux operating system compatible with the Red Hat Linux. The NEC support portal below lists the operating systems and their kernel versions verified for the SX-Aurora TSUBASA.

[SX-Aurora TSUBASA]Verified Linux kernel

https://www.support.nec.co.jp/en/View.aspx?id=4140100078

As listed in the above page, only updated kernels are supported and they are not included in the ISO image of each distribution. So it is necessary to update the kernel before use. Also, to avoid kernel update to a version that is not verified, please

configure the yum command using the file /etc/yum.conf so that kernel packages are not updated. Please refer to 2.2 for the configuration.

1.2.3 Related URL

See also the following site about SX-Aurora TSUBASA.

✓ NEC Aurora Forum(https://sxauroratsubasa.sakura.ne.jp/)

1.3 Examples of System Configuration

This section illustrates system configuration examples of the SX-Aurora TSUBASA.

Configuration 1: Standalone (Single Vector Island (VI))

Figure 1 illustrates the SX-Aurora TSUBASA software to be installed on the Vector Host (VH).

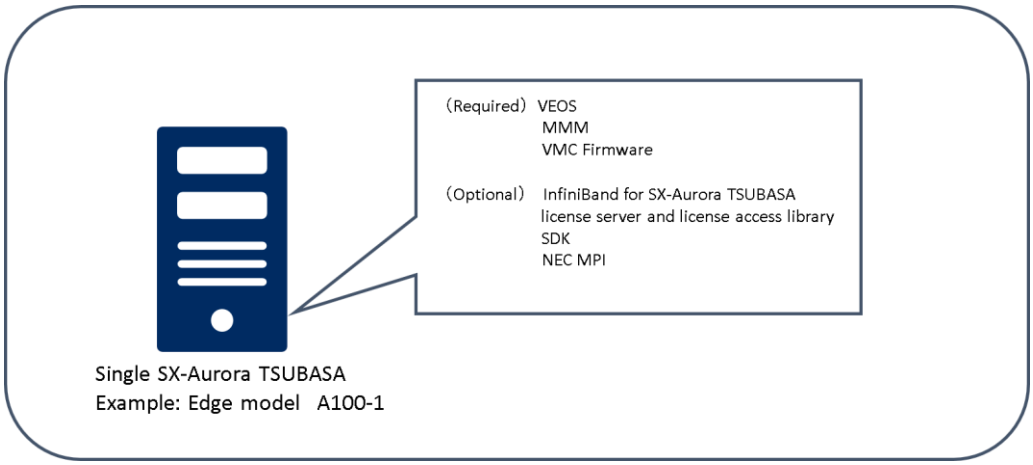

Figure 1 Configuration 1: Standalone (Single VI)

Configuration 2: Multiple VIs, a Management Server, and a Frontend Machine

In this case, software license management is performed on the management server and programs can be compiled on the frontend machine.

Figure 2 illustrates the SX-Aurora TSUBASA software to be installed on the VHs, management server, and frontend machine.

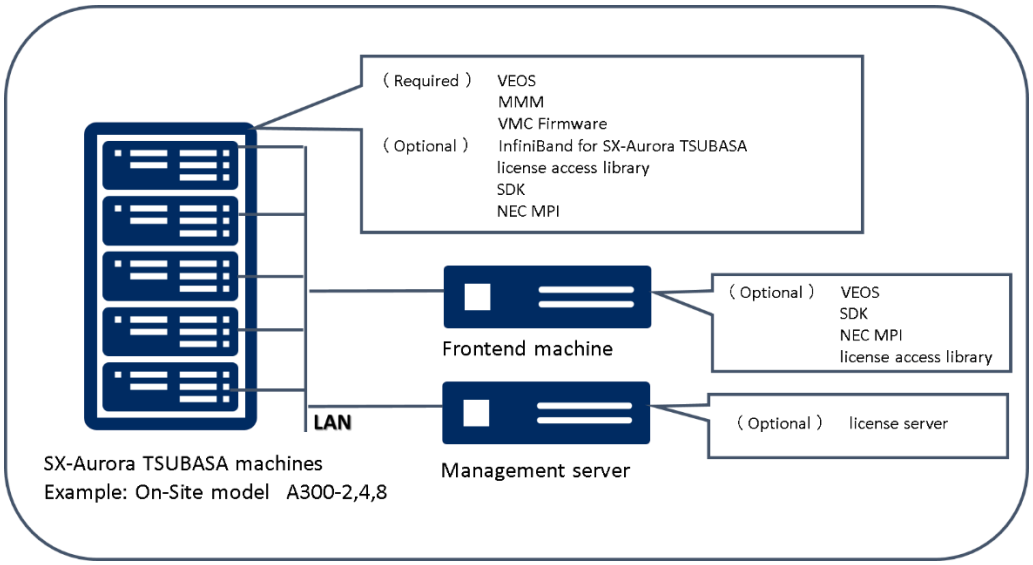

Figure 2 Configuration 2: Multiple VIs, a Management server, and a Frontend Machine

Configuration 3: Large Scale System

Please contact our sales or SE.

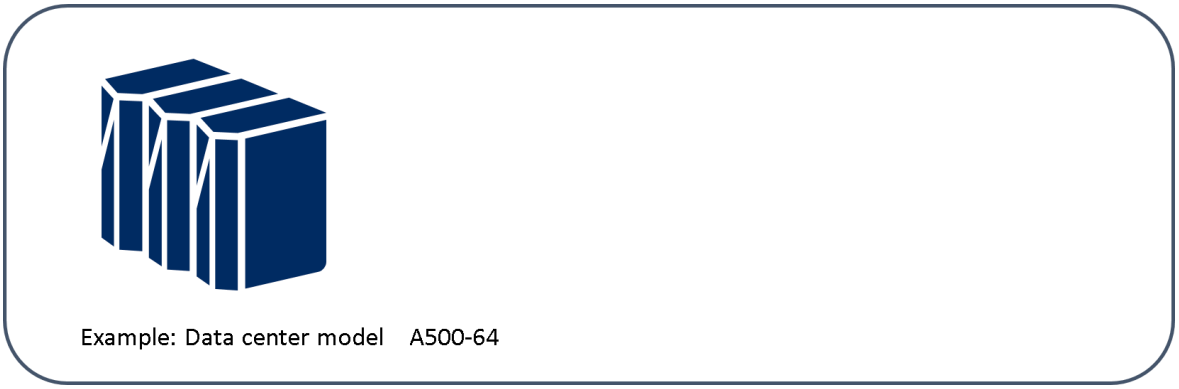

Figure 3 Configuration 3: Large Scale System

1.4 Difference of software installation

Installation of software is different depending on wheather you purchased the paid software (1.1.2) and have the PP support contract or not and wheather you can direct access NEC yum repository or not.

1.4.1 With the paid software and PP Support contract

When you purchased the paid software and have the PP support contract, you can install the software packages from the NEC yum repository with restricted access. The serial number of the support pack is required to access the yum repository with restricted access. Please refer to Section 3.1.2 for the serial number.

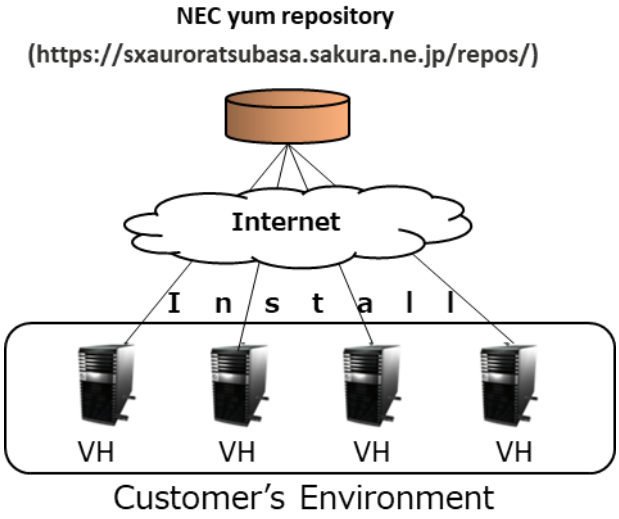

Figure 4 Installation from the yum repository

If your SX-Aurora TSUBASA system does not have direct access to the Internet, you can install them by setting up the yum repository in the local environment. Please refer to 3.1.3 for how to set up the local yum repository.

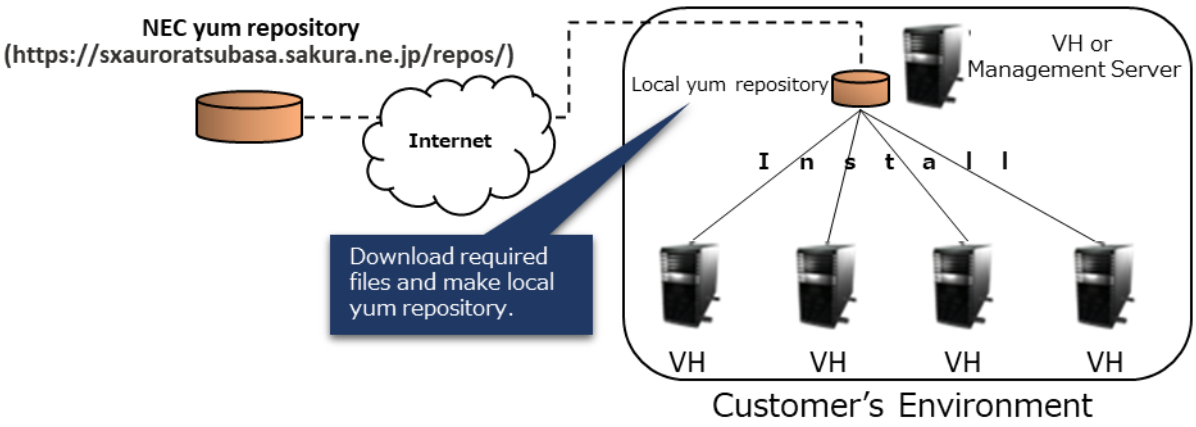

Figure 5 Installation from the local yum repository

1.4.2 Without the paid software or PP Support contract

If you purchased the paid software and do not have the PP support contract, you can install the paid software packages from the internet delivery product download service.

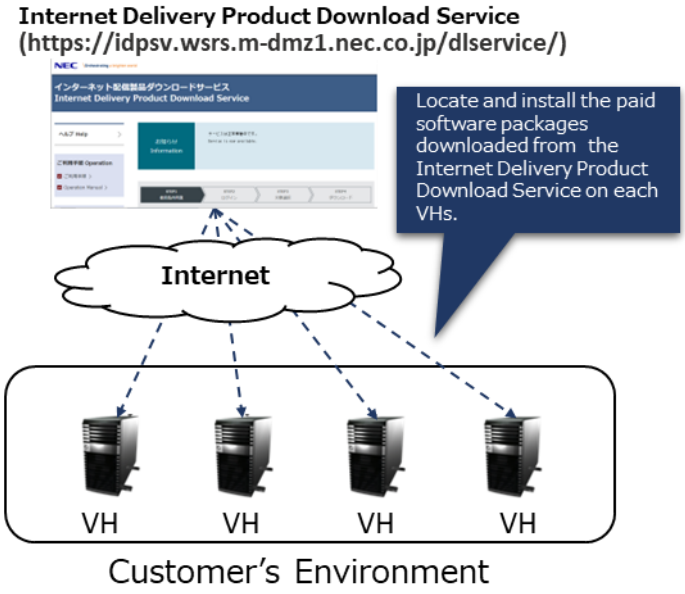

Figure 6 Installation of packages downloaded in the Internet Delivery Product Download Service

As for the free software packages, you can install them from the NEC yum repository with free access. If your SX-Aurora TSUBASA system does not have direct access to the Internet, you can install them by setting up the yum repository in the local environment. Please refer to 3.1.3 for how to set up the local yum repository.

1.5 Software Installation procedure

You can install the SX-Aurora TSUBASA software onto VHs according to Figure 7. Please refer to Chapter2 and 2.3 for details. Also, configure the software referring to Chapter4 after the installation.

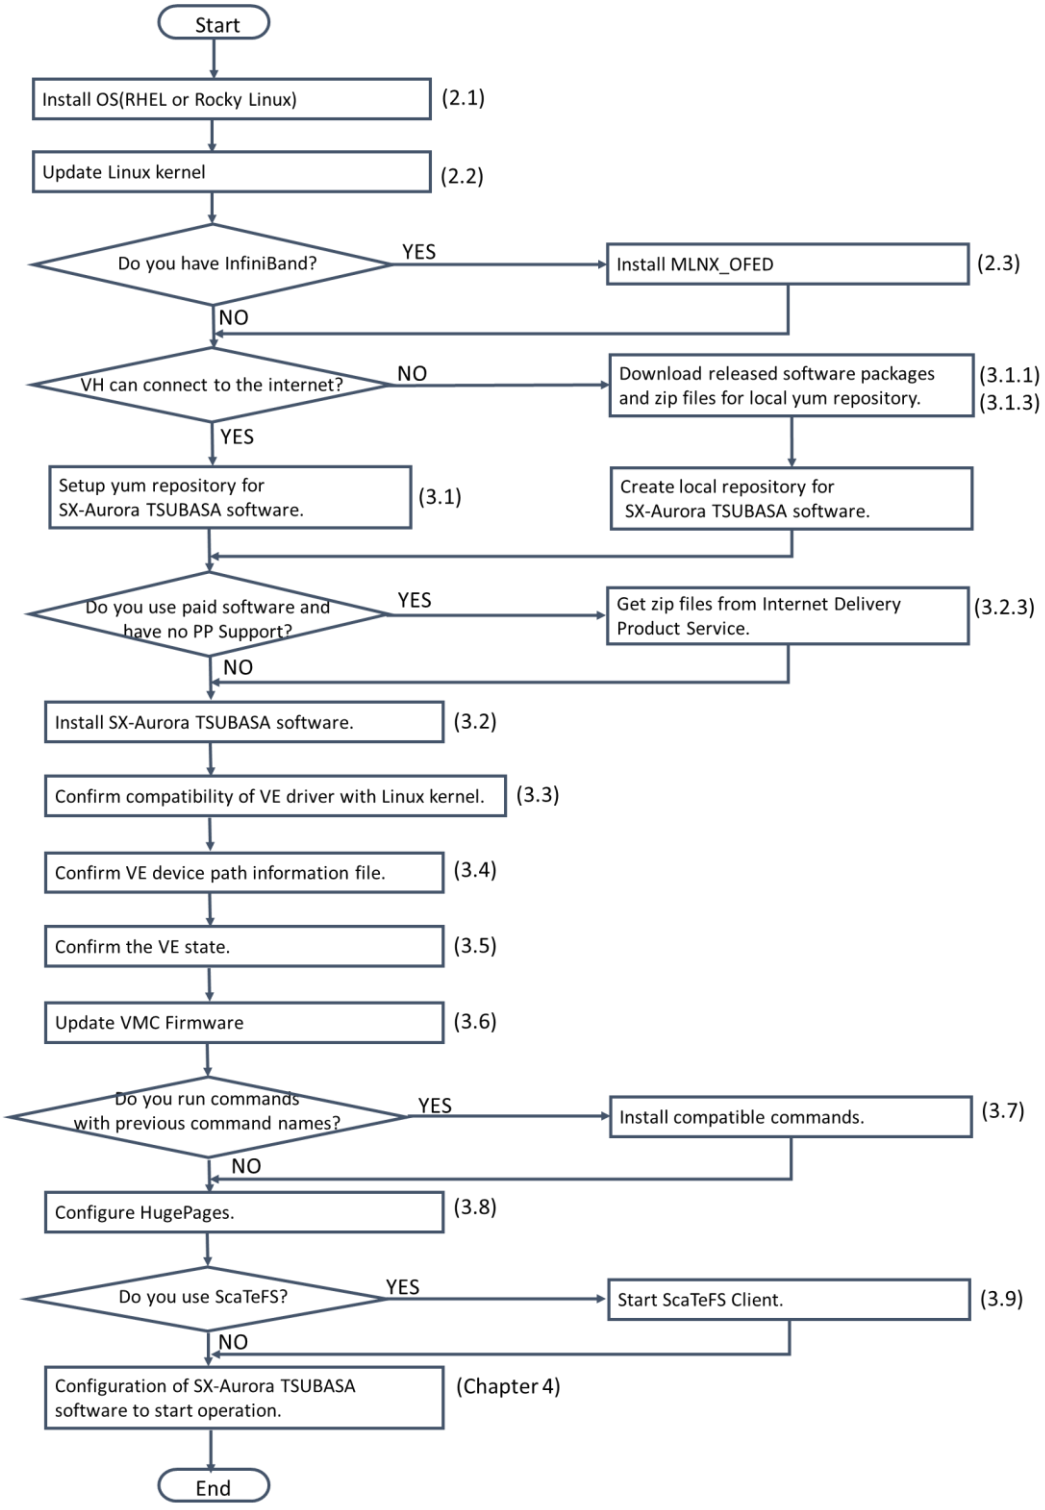

Figure 7 Installation procedure

Chapter2 Installation OS and the related Software

This chapter explains the procedure installing OS and the related software.

2.1 OS Installation onto the VHs

2.1.1 Linux OS installation

Before installing the SX-Aurora TSUBASA software, install a supported operating system on the VHs with reference to the following information.

- [SX-Aurora TSUBASA]Verified Linux kernel
  https://www.support.nec.co.jp/en/View.aspx?id=4140100078
- Installation of the operating system
  - Red Hat Enterprise Linux
    The Red Hat Customer Portal
    ✓ Product Documentation for Red Hat Enterprise Linux 8 - Installing
  - Rocky Linux
    The documentation on the Rocky Linux site

2.1.2 Linux OS yum repository

Set up the yum repository for the Linux OS so that additional packages for the OS can be installed by the yum command. This is required to install SX-Aurora TSUBASA software. There are two ways to set up the yum repository. One is the way to use the Linux OS installation media, and the other is the way to use official repository on the internet.

This section explains about the settings to use yum repository which accesses to Rocky Linux 8 OS installation media (DVD). In the case of other OS, he following repository name, file name and GPG key file should be changed according to the target OS version.

(1) Mount OS installation media (DVD).

Put the OS installation DVD into the DVD drive and mount it at an appropriate directory. In the following example, /media/cdrom is used.

```
# mkdir /media/cdrom
```

```
# mount /dev/cdrom /media/cdrom
```

(2) Configure repository settings

To enable to install software from the OS installation media (DVD) by the yum command, save the original repository configuration files and create a new configuration file for the OS installation DVD.

First, save the original configuration files under /etc/yum.repos.d.

```
# cd /etc/yum.repos.d
# mkdir Rocky-repos.d
# mv Rocky-* Rocky-repos.d
# cp Rocky-repos.d/Rocky-Media.repo .
```

Edit /etc/yum.repos.d/Rocky-Media.repo as follwing:

```
[media-baseos]
name=Rocky Linux $releasever - Media - BaseOS
baseurl=file:///media/cdrom/BaseOS
gpgcheck=1
enabled=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-rockyofficial

[media-appstream]
name=Rocky Linux $releasever - Media - AppStream
baseurl=file:///media/cdrom/AppStream
gpgcheck=1
enabled=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-rockyofficial
```

Please keep the DVD mounted until installation of the SX-Aurora TSUBASA software is completed.

2.2 Kernel Update

Please update the kernel on the VHs to a version verified for the SX-Aurora TSUBASA, and reboot them.

The NEC support portal below lists the operating systems and their kernel versions verified for the SX-Aurora TSUBASA.

[SX-Aurora TSUBASA]Verified Linux kernel

https://www.support.nec.co.jp/en/View.aspx?id=4140100078

After the update, to avoid kernel update to a version that is not verified, please configure the yum command using the file /etc/yum.conf so that kernel packages are not updated. The following is an example of the description in the file /etc/yum.conf to avoid kernel update, where 'exclude=kernel*' is specified.

```
# vi /etc/yum.conf
[main]
exclude=kernel*
```

You can install Rocky Linux 8.10 verified kernel(kernel-4.18.0-553.8.1.el8_10.x86_64) as following procedure.

(1) Setting of Rocky Linux 8.10 repository

Create /etc/yum.repos.d/Rocky-BaseOS.repo as below to connect to http://dl.rockylinux.org/pub/rocky/8.10/. Edit the file, and comment out "mirrorlist=https://mirrors.rockylinux.org/mirrorlist?arch=$basearch&repo=BaseOS-$releasever",
and add
"baseurl=http://dl.rockylinux.org/pub/rocky/8.10/BaseOS/$basearch/os/".

```
# cd /etc/yum.repos.d
# cp Rocky-repos.d/Rocky-BaseOS.repo .
# vi Rocky-BaseOS.repo
# diff Rocky-repos.d/Rocky-BaseOS.repo Rocky-BaseOS.repo
13c13
<
mirrorlist=https://mirrors.rockylinux.org/mirrorlist?arch=$basearch&repo=BaseOS-
$releasever
---
>
#mirrorlist=https://mirrors.rockylinux.org/mirrorlist?arch=$basearch&repo=BaseOS-
$releasever
14a15
> baseurl=http://dl.rockylinux.org/pub/rocky/8.10/BaseOS/$basearch/os/
```

(2) Installing the kernel (kernel-4.18.0-553.8.1.el8_10.x86_64)

```
#  dnf  --disableexcludes=all  install  kernel-4.18.0-553.8.1.el8_10  kernel-headers-
4.18.0-553.8.1.el8_10
Last metadata expiration check: 0:03:09 ago on Wed Sep 25 10:01:15 2024.
Dependencies resolved.
================================================================================
 Package           Arch       Version                   Repository   Size
================================================================================
```

```
Installing:
 kernel             x86_64      4.18.0-553.8.1.el8_10       baseos      10 M
 kernel-headers      x86_64      4.18.0-553.8.1.el8_10       baseos      12 M
Installing dependencies:
 kernel-core        x86_64      4.18.0-553.8.1.el8_10       baseos      43 M
 kernel-modules      x86_64      4.18.0-553.8.1.el8_10       baseos      36 M

Transaction Summary
================================================================================
Install  4 Packages

Total download size: 101 M
Installed size: 102 M
Is this ok [y/N]: y
   :
Installed:
  kernel-4.18.0-553.8.1.el8_10.x86_64
  kernel-core-4.18.0-553.8.1.el8_10.x86_64
  kernel-headers-4.18.0-553.8.1.el8_10.x86_64
  kernel-modules-4.18.0-553.8.1.el8_10.x86_64

Complete!

# reboot
```

⚠ Note

Regarding the kernel version of the host machine, we recommend that you update the kernel to the latest version among the kernels that have been verified for operation in principle. If the latest version is not used, ve_drv module in ve_drv-kmod package may not be loaded and VE may not be ONLINE after installing Aurora software. In that case, you will need to downgrade vp-kmod and ve_drv-kmod packages. Please refer to "3.3 Confirmation of VE Driver Compatible with Linux kernel" for details.

2.3 Installation of MLNX_OFED (Optional)

If you use InfiniBand with the SX-Aurora TSUBASA, install MLNX_OFED onto the VHs according to Table 4.

Table 4 Correspondence between the OS Versions and MLNX_OFED

OS	MLNX_OFED	Supported VE	
		VE1/VE2	VE3
RHEL/Rocky Linux 8.10	MLNX_OFED 23.10-3.2.2.0-LTS	✓	✓

You can obtain the ISO image of MLNX_OFED and related documents from the NVIDIA official home page.

The NVIDIA Official Home Page

https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed

⚠ Note

1. When installing MLNX_OFED on C401-8 or a model with ConnectX-7 attached, please install it with the "--without-fw-update" option. If you install MLNX_OFED 23.10-3.2.2.0-LTS without specifying this option, HCA-FW version 28.39.3560 will be automatically installed. If HCA-FW version 28.39.3560 is installed, please downgrade the HCA-FW after the reboot. (Please refer to 2.3.1)
2. Install software packages required for MLNX_OFED referring to "Hardware and Software Requirements" in the release notes of MLNX_OFED before the installation. If missing packages are displayed while installation of MLNX_OFED, please install those packages. MLNX_OFED documents are available in the following URL.
   https://docs.nvidia.com/networking/software/adapter-software/index.html
3. If the Linux kernel is updated, installation with the mlnxofedinstall command may fail on compatibility check. In this case, please execute the mlnx_add_kernel_support.sh --kmp command or mlnxofedinstall --add-kernel-support --kmp command.
4. Package kernel-rpm-macros may be required to execute mlnxofedinstall --add-kernel-support --kmp. In this case, please comment out exclude=kernel* in /etc/yum.conf temporarily and install kernel-rpm-macros.
5. If the software which depends on InfiniBand has been installed, installation may fail with following messages. In this case, please uninstall displayed packages and retry installation.

```
Error: One or more packages depends on MLNX_OFED_LINUX.
These packages should be removed before uninstalling MLNX_OFED_LINUX:
```

After the installation, please reload the InfiniBand driver or reboot the VH.

```
# /etc/init.d/openibd restart
```

⚠ Note

1. If the firmware is updated, reboot is mandatory. Please refer to the MLNX_OFED documents for details.
2. If reloading InfiniBand driver is failed, please reboot the VH.
3. After the reboot, the status of MLNX_OFED can be checked by the following command.

```
# systemctl status openibd
```

2.3.1 Notice when using model B403-8, C401-8 or a model with ConnectX-7 for IB-HCAs attached

When using model B403-8, C401-8 or a model with ConnectX-7 for IB-HCA attached,please use HCA-FW version 28.39.1002 for IB HCA (NDR200 compatible ConnectX-7).

1. Check the version of HCA-FW

   Check if the HCA-FW version is 28.39.1002 or not.

```
# ibstat | grep -A2 MT4129
       CA type: MT4129
       Number of ports: 1
       Firmware version: 28.39.1002
--
       CA type: MT4129
       Number of ports: 1
       Firmware version: 28.39.1002
```

2. Apply HCA-FW

   If the HCA-FW version is not 28.39.1002, download Firmware version 28.39.1002 from the NVIDIA site below and install the HCA-FW.

   https://network.nvidia.com/support/firmware/connectx7

- How to download

  Click the ARCHIVE VERSION tab

  -> Select Version 28.39.1002, OPN MCX75310AAS-HEA, PSID MT_0000000844

  -> Select Download fw-ConnectX7-rel-28_39_1002-MCX75310AAS-HEA_Ax-UEFI-

14.32.12-FlexBoot-3.7.201.signed to download the HCA-FW

After download, please extract the zip file.

- How to update the HCA-FW

  If the displayed firmware version is older than 28.39.1002, please execute the following as root.

```
# mst start
# mlxfwmanager -d /dev/mst/mt4129_pciconf0 -i fw-ConnectX7-rel-28_39_1002-MCX75310AAS-HEA_Ax-UEFI-14.32.12-FlexBoot-3.7.201.signed.bin -u -y
# mlxfwmanager -d /dev/mst/mt4129_pciconf1 -i fw-ConnectX7-rel-28_39_1002-MCX75310AAS-HEA_Ax-UEFI-14.32.12-FlexBoot-3.7.201.signed.bin -u -y
```

  After applying, please restart.

```
# reboot
```

- How to downgrade the HCA-FW

  If the displayed firmware version is newer than 28.39.1002, please execute the following as root.

```
# mst start
# mlxfwmanager -d /dev/mst/mt4129_pciconf0 -i fw-ConnectX7-rel-28_39_1002-MCX75310AAS-HEA_Ax-UEFI-14.32.12-FlexBoot-3.7.201.signed.bin -u -y -f
# mlxfwmanager -d /dev/mst/mt4129_pciconf1 -i fw-ConnectX7-rel-28_39_1002-MCX75310AAS-HEA_Ax-UEFI-14.32.12-FlexBoot-3.7.201.signed.bin -u -y -f
```

  After applying, please restart.

```
# reboot
```

Chapter3 Installation the SX-Aurora TSUBASA software

3.1 Prepration of installation

Set up the yum repository to install the SX-Aurora TSUBASA software.

You can use the NEC yum repository on the Internet or the one locally set up.

⚠ Note

For the purpose of the stable securing of operation and prevention of unauthorized use of yum repository service, we record your IP address and access-ID when you access the yum repository on the Internet.

3.1.1 The files to use yum repository

It is necessary to install SX-Aurora TSUBASA software repository setting files, in the case from not only Internet repository but also your local repository. The architecture configuration file contained in this rpm file determines the default architecture. The information in this configuration file is used by the NEC compiler to generate binaries for the VE1 or VE3 architecture when it compiles.

- To set the default architecture to VE1

  https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve1-3.0-1.noarch.rpm

- To set the default architecture to VE3

  https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve3-3.0-1.noarch.rpm

This file includes the following files:

File	Description
/etc/opt/nec/ve/default.conf	Configuration file for default VE architecture
/etc/pki/rpm-gpg/RPM-GPG-KEY-TSUBASA-soft	GPG public Key
/etc/yum.repos.d/TSUBASA-restricted.repo	Configuration file for the yum repository with restricted access
/etc/yum.repos.d/TSUBASA-repo.repo	Configuration file for the yum repository with free access

File	Description
/opt/nec/ve/sbin/TSUBASA-groups-remark.sh	Script to update the group status
/opt/nec/ve/sbin/terminate-all-ve-services	Script to stop services
/opt/nec/ve/sbin/start-all-ve-services	Script to restart services
/opt/nec/ve/sbin/setup-ve-infiniband.sh	Script to configure InfiniBand

The procedure to install the rpm is explained in 3.2.

3.1.2 The ID to access the restricted repository

To install the paid software, you need to access the yum repository with restricted access and you need the support pack or support contract to access it.

Edit the yum configuration file /etc/yum.repos.d/TSUBASA-restricted.repo for the paid software based on the contract details you have: Enter the 16 digits of the serial number on the serial number card, excluding the hyphens, in the username field, the right eight digits in the password field, and 1 in the enabled field.
If you have purchased the support pack, the serial number card is included in the file you can download from the internet delivery product download service. If you have the support contract, please contact our sales. Figure 8 Serial Number Card shows an example of the serial number card.

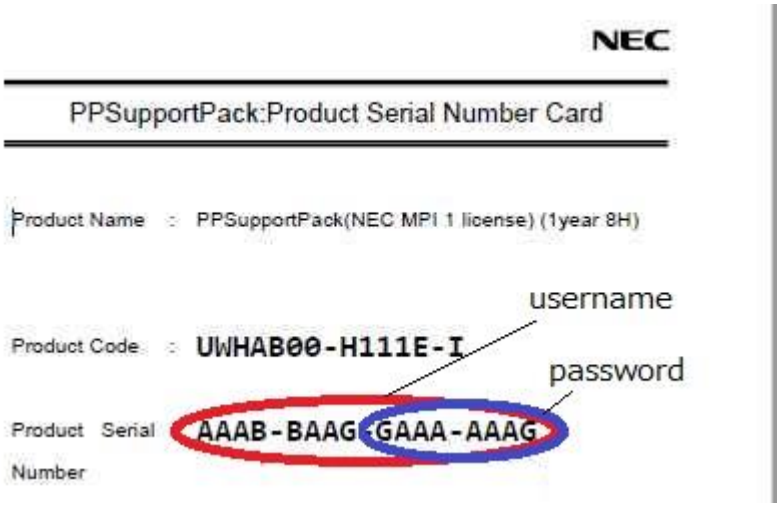

Figure 8 Serial Number Card

- Accessing to NEC yum repository

The following example shows the description in the yum configuration file to enable the yum repository with restricted access.

```
# vi /etc/yum.repos.d/TSUBASA-restricted.repo
[nec-sdk]
  …
username=<serial number for NEC SDK>
password=<the right eight digits of the serial number for NEC SDK>
enabled=1    ←update 0 to 1
  …
[nec-mpi]
  …
username=<serial number for NEC SDK *>
password=<the right eight digits of the serial number for NEC SDK *>
enabled=1    ←update 0 to 1
  …
```

* If you purchased the PPSP for NEC SDK after January 2021, set the PPSP for NEC SDK serial No to nec-mpi section. If you purchased the PPSP for NEC SDK before December 2020, set the PPSP for NEC MPI serial No to nec-mpi section.

* If you have a contract for PP Support but have not purchased SDK-SE or Compat C++, please enter 0 in the enabled field for nec-sdk.

- Not accessing to NEC yum repository

When not accessing to NEC yum repository directly from your VH, please authenticate your ID to download the zip files for local yum repository.

3.1.3 In the case of not accessing to NEC repository from VH

If your SX-Aurora TSUBASA system does not have direct access to the Internet, it is necessary to set up a locally accessible yum repository.

(1) Zip files to create local yum repository

You can obtain the zip files of the yum repository from the following links.

The x of RHEL/Rocky Linux 8.x is your using RHEL/Rocky Linux version.

Under https://sxauroratsubasa.sakura.ne.jp/repos/

	RHEL/Rocky Linux 8.x
Yum repository, configuration files, etc.	TSUBASA-soft-release-ve1-3.0-1.noarch.rpm or TSUBASA-soft-release-ve3-3.0-1.noarch.rpm

	RHEL/Rocky Linux 8.x
VEOS, MMM, etc.	TSUBASA-repo_el8.x.zip
SDK-CE	community/sdk/sdk_el8.zip
	community/mpi/mpi_mofed5_el8.zip
SDK-SE, Compat C++	restricted/sdk/sdk_el8.zip
	restricted/mpi/mpi_mofed5_el8.zip
SDK Runtime	runtime/sdk/sdk_el8.zip
	runtime/mpi/mpi_mofed5_el8.zip
NQSV	restricted/nqsv/nqsv_el8.zip
ScaTeFS	restricted/scatefs/scatefs_el8.x.zip

zip files under community and restricted which is used to install SDK-CE, SDK-SE and Compat C++ includes contents in zip files under runtime. Therefore, when you download the zip files from community or restricted, you do not have to download the runtime one.

Please note that access to the files for the repository with restricted access requires the username (the 16 digits of the serial number) and password (the right eight digits of the serial number) of the PP support contract.

Download the following release package, which includes the configuration files for the yum repository. Please install the package because it will change the yum repository configuration.

- To set the default architecture to VE1

  TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
- To set the default architecture to VE3

  TSUBASA-soft-release-ve3-3.0-1.noarch.rpm

When you use VEOS, MMM, etc. or ScaTeFS, download the following release packages for your OS version.

	OS version	File link
VEOS, MMM, etc.	8.10	TSUBASA-repo_el8.10.zip
ScaTeFS	8.10	restricted/scatefs/scatefs_el8.10.zip

When you use SDK or NQSV on RHEL/Rocky Linux 8.x, download the following release packages for software.

		File link (RHEL/Rocky Linux 8.x)
SDK (except MPI)	SDK-CE	community/sdk/sdk_el8.zip
	SDK-SE, Compat C++	restricted/sdk/sdk_el8.zip
	SDK Runtime	runtime/sdk/sdk_el8.zip
SDK (MPI)	SDK-CE	community/mpi/mpi_mofed5_el8.zip
	SDK-SE, Compat C++	restricted/mpi/mpi_mofed5_el8.zip
	SDK Runtime	runtime/mpi/mpi_mofed5_el8.zip
NQSV	NQSV	restricted/nqsv/nqsv_el8.zip

(2) Setup of the local yum repository

If your SX-Aurora TSUBASA system is not standalone, set up the yum repository for the SX-Aurora TSUBASA software on one of the VHs or the management server. Otherwise, set it up on the VH. Hereafter, the host where you set up the yum repository is called the repository server. The following describes how to set it up on the repository server.

1. If your SX-Aurora TSUBASA system is not standalone, install and run the Apache HTTP server on the repository server. Otherwise, the Apache HTTP server is not needed.

```
# yum -y install httpd
# systemctl start httpd
# systemctl enable httpd
```

2. If you SX-Aurora TSUBASA system is not standalone, create a directory for the yum repository under the HTTP server document path (/var/www/html). Otherwise, create a directory for the yum repository under an arbitrary directory.

```
# mkdir -p /path/to/repos
```

3. Place all the files you have downloaded under the directory and install the release package.

- To set the default architecture to VE1

```
# cd /path/to/repos
# yum install ./TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

- To set the default architecture to VE3

```
# cd /path/to/repos
# yum install ./TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

⚠ Note

When the following messages are displayed during installation, it indicates that a configuration file was saved in a different file name (added .rpmnew extension) so that the new file doesn't overwrite an existing configuration. The new file includes configurations for new features. If user doesn't use new features, there is nothing to do about *.rpmnew file. If user will use new features, user needs to copy newly added configurations to the existing file and modify them for user's environment.

```
warning: /etc/yum.repos.d/TSUBASA-repo.repo created as /etc/yum.repos.d/TSUBASA-re
po.repo.rpmnew
warning: /etc/yum.repos.d/TSUBASA-restricted.repo created as /etc/yum.repos.d/TSUB
ASA-restricted.repo.rpmnew
```

4. Specify the location of the local yum repository in the "baseurl=" line in the configuration files for the repository with free access (/etc/yum.repos.d/TSUBASA-repo.repo) and the repository with restricted access (/etc/yum.repos.d/TSUBASA-restricted.repo), and enter 1 in the enabled field.

- Example for VEOS, MMM, etc.

(Before)

```
baseurl=https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-repo_el8.10
```

(After)

```
baseurl=file:///path/to/repos/TSUBASA-repo_el8.10
```

- Example for SDK-CE

(Before)

```
baseurl=https://sxauroratsubasa.sakura.ne.jp/repos/community/sdk/sdk_el8
```

(After for standalone systems)

```
baseurl=file:///path/to/repos/sdk_el8
```

(After for non-standalone systems)

It is not necessary to describe the path specified by DocumentRoot(e.g. /var/www/html).

```
Baseurl=http://address.or.name.of.server/path/to/repos/sdk_el8
```

Delete the yum cache

```
# yum clean all
```

5. Expand the zip files to get the repository

```
# cd /path/to/repos/
# find . -name '*.zip' | xargs -n1 unzip -o
# rm *.zip
```

(3) Update local yum repository

Place and expand the zip files downloaded in step (1) in the directory for the yum repository (/path/to/repos in the example below) on the repository server.

```
# cd /path/to/repos/
# find . -name '*.zip' | xargs -n1 unzip -o
# rm *.zip
```

When you have updated the OS on the VHs, it is also necessary to update the version number in the "baseurl=" line in the configuration files to the one corresponding to the OS version.

- Example for VEOS, MMM, etc.

(Before)

```
baseurl=file:///path/to/repos/TSUBASA-repo_el7.6
```

(After)

```
baseurl=file:///path/to/repos/TSUBASA-repo_el8.10
```

- Example for SDK

(Before)

```
baseurl=file:///path/to/repos/sdk_el7
```

(After)

```
baseurl=file:///path/to/repos/sdk_el8
```

Delete the yum cache.

```
# yum clean all
```

3.2 Software Installation

Install the SX-Aurora TSUBASA software with the value of the shell variable TSUBASA_GROUPS set to the group names to be installed according to whether the SX-Aurora TSUBASA has InfiniBand and which paid software you have purchased. Please refer to 3.2.4 for the available group names.

⚠ Note

When you cannot access to the yum repository, try it again after few minutes. The service may not be available due to server maintenance, so please check the following as well.

https://sxauroratsubasa.sakura.ne.jp/forums/

3.2.1 With the paid software and PP Support - direct access to NEC yum repository

- Install SX-Aurora TSUBASA software repository setting files
  - To set the default architecture to VE1

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

  - To set the default architecture to VE3

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

⚠ Note

When the following messages are displayed during installation, it indicates that a configuration file was saved in a different file name (added .rpmnew extension) so that the new file doesn't overwrite an existing configuration. The new file includes configurations for new features. If user doesn't use new features, there is nothing to do about *.rpmnew file. If user will use new features, user needs to copy newly added configurations to the existing file and modify them for user's environment.

```
warning: /etc/yum.repos.d/TSUBASA-repo.repo created as /etc/yum.repos.d/TSUBASA-repo.repo.rpmnew
warning: /etc/yum.repos.d/TSUBASA-restricted.repo created as /etc/yum.repos.d/TSUBASA-restricted.repo.rpmnew
```

- Install SX-Aurora Software

  for the environment to compile and execute programs

```
# TSUBASA_GROUPS="ve-devel ve-infiniband nec-sdk-devel nec-mpi-devel nqsv-execution scatefs-client-tsubasa"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

  for the environment only to execute programs

```
# TSUBASA_GROUPS="ve-runtime ve-infiniband nec-sdk-runtime nec-mpi-runtime nqsv-execution scatefs-client-tsubasa"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

  for the frontend machine

```
# TSUBASA_GROUPS="ve-frontend nec-sdk-frontend nec-mpi-frontend"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

3.2.2 With the paid software and PP Support - local yum repository

- Setup of the local yum repository

  If your SX-Aurora TSUBASA system does not have direct access to the Internet, it is necessary to set up a locally accessible yum repository according to 3.1.3.

- Install SX-Aurora TSUBASA software repository setting files

Run the following commands, and setup the files to access to local yum repository.

- To set the default architecture to VE1

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

- To set the default architecture to VE3

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

⚠ Note

When the following messages are displayed during installation, it indicates that a configuration file was saved in a different file name (added .rpmnew extension) so that the new file doesn't overwrite an existing configuration. The new file includes configurations for new features. If user doesn't use new features, there is nothing to do about *.rpmnew file. If user will use new features, user needs to copy newly added configurations to the existing file and modify them for user's environment.

```
warning: /etc/yum.repos.d/TSUBASA-repo.repo created as /etc/yum.repos.d/TSUBASA-repo.repo.rpmnew
warning: /etc/yum.repos.d/TSUBASA-restricted.repo created as /etc/yum.repos.d/TSUBASA-restricted.repo.rpmnew
```

- Install SX-Aurora Software

for the environment to compile and execute programs

```
# TSUBASA_GROUPS="ve-devel ve-infiniband nec-sdk-devel nec-mpi-devel nqsv-execution scatefs-client-tsubasa"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

for the environment only to execute programs

```
# TSUBASA_GROUPS="ve-runtime ve-infiniband nec-sdk-runtime nec-mpi-runtime nqsv-execution scatefs-client-tsubasa"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

for the frontend machine

```
# TSUBASA_GROUPS="ve-frontend nec-sdk-frontend nec-mpi-frontend"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

3.2.3 Without the paid software or PP Support

- Install free software packages

1. Setup yum repository

✧ From NEC yum repository with free access

- To set the default architecture to VE1

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

- To set the default architecture to VE3

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

✧ From Local yum repository

Run the following commands, and setup the files to access to local yum repository.

- To set the default architecture to VE1

```
# cd /path/to/repos
# yum install ./TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

- To set the default architecture to VE3

```
# cd /path/to/repo
# yum install ./TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

⚠ Note

When the following messages are displayed during installation, it indicates that a configuration file was saved in a different file name (added .rpmnew extension) so that the new file doesn't overwrite an existing configuration. The new file includes configurations for new features. If user doesn't use new features, there is nothing to do about *.rpmnew file. If user will use new features, user needs to copy newly added configurations to the existing file and modify them for user's environment.

```
warning: /etc/yum.repos.d/TSUBASA-repo.repo created as /etc/yum.repos.d/TSUBASA-repo.repo.rpmnew
warning: /etc/yum.repos.d/TSUBASA-restricted.repo created as /etc/yum.repos.d/TSUBASA-restricted.repo.rpmnew
```

2. Prepare the paid software packages

Download the latest zip files of the paid software packages in the internet delivery

product download service. Then transfer the files onto the VHs. If you have any old zip files of the paid software packages, delete them and then unzip the latest zip files.

```
# find . -name '*.zip' | xargs -n1 unzip -o
```

3. Install free software packages

✧ for the environment to compile and execute programs

```
# yum install devel/*.rpm
# TSUBASA_GROUPS="ve-devel ve-infiniband"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

✧ for the environment only to execute programs

```
# yum install runtime/*.rpm
# TSUBASA_GROUPS="ve-runtime ve-infiniband"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

✧ for the frontend machine

```
# yum install frontend/*.rpm
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh ve-frontend
# yum group install ve-frontend
```

✧ when using NQSV

```
 # yum install NQSV-client/*.rpm
```

✧ when using ScaTeFS/Client

If you use the ScaTeFS client, install the packages corresponding to your environment as follows after installation of the ve-infiniband package group or ve-infiniband-mofedX.X package group.

・ When installing on RHEL/Rocky Linux 8.10

```
# yum install scatefs-client/el8.10/scatefs-client-tsubasa/*.rpm
```

3.2.4 Package groups for TSUBASA_Software

The package groups for various types of machines are shown in the following tables. These package groups can be specified using the shell variable TSUBASA_GROUPS.

Table 5 Package Groups for SX-Aurora TSUBASA

Package group	Description	VH for compiling and executing program	VH for executing program	Frontend machine
ve-devel	VE Application	✓	-	-
ve-runtime		-	✓	-
ve-frontend		-	-	✓
ve-infiniband	InfiniBand for SX-Aurora TSUBASA	✓	✓	-
nec-sdk-devel	SDK (except MPI)	✓	-	-
nec-sdk-frontend		-	-	✓
nec-sdk-runtime		-	✓	-
	SDK Runtime (except MPI)			
nec-mpi-devel	SDK (MPI)	✓	-	-
nec-mpi-frontend		-	-	✓
nec-mpi-runtime		-	✓	-
	SDK Runtime (MPI)			
nqsv-execution	NQSV	✓	✓	-
scatefs-client-tsubasa	ScaTeFS Client	✓	✓	-

✓ The packages for the each Package group are described in Appendix B.

Table 6 Package Group/NQSV

Package	Group Name	
	nqsv-batch	nqsv-execution
NQSV-API	✓	-
NQSV-Client	✓	✓
NQSV-JobManipulator	✓	-
NQSV-JobServer	-	✓
NQSV-ResourceManager	✓	-

Table 7 Package Group/ScaTeFS Client

Package	Group Name		
	scatefs-client-tsubasa	scatefs-client-tsubasa-mofedX.X (*1)	scatefs-client-tsubasa-container
kmod-scatefs-client-modules-mofedX.X	✓	✓	-
scatefs-client-libscatefsib_ve	✓	✓	✓
scatefs-client-mount-utils	✓	✓	-
scatefs-client-rcli-utils	✓	✓	-
scatefs-client-sys-utils	✓	✓	-
scatefs-client-usr-utils	✓	✓	✓

(*1) Use this package group when using a ScaTeFS Client that supports specific version of MLNX_OFED. Use the "scatefs-client-tsubasa" group instead of the "scatefs-client-tsubasa-mofedX.X" group unless there is a specific reason.

Table 8 Package Group/ScaTeFS Server

Package	Group Name	
	scatefs-server	scatefs-server-monitoring (*1)
scatefs-srv	-	✓
scatefs-mon	✓	-

(*1) The group "scatefs-server-monitoring" is supported on RHEL 7.7 or later.

3.3 Confirmation of VE Driver Compatible with Linux kernel

Downgrading VE driver (vp-kmod and ve_drv-kmod packages) is required when the installed VE driver is not compatible with Linux kernel in host machine.

Regarding the kernel version of the host machine, we recommend that you update the kernel to the latest version among the kernels that have been verified for operation in principle (2.2 Kernel Update). If the latest version is not used, ve_drv module in ve_drv-kmod package may not be loaded and VE may not be ONLINE after installing Aurora software. In that case, you will need to downgrade vp-kmod and ve_drv-kmod packages.

3.3.1 Downgrade Confirmation

Check to see if downgrades are needed. First, check to see if the ve_drv module is loaded.

```
# lsmod | grep ve_drv
ve_drv              262144  128
vp                  20480  1 ve_drv
```

If ve_drv is displayed, ve_drv is loaded. Please go to section 3.4.

```
# lsmod | grep ve_drv
(Nothing displayed)
```

If nothing is displayed, it has not been loaded. Check the Linux kernel version of VH with the following command.

```
# uname -r
4.18.0-553.8.1.el8_10.x86_64
```

Next, check the versions of the vp-kmod and ve_drv-kmod packages and make sure that the coreesponding packages are installed for your kernel.

[SX-Aurora TSUBASA]Verified Linux kernel

https://www.support.nec.co.jp/en/View.aspx?id=4140100078

```
# rpm -q vp-kmod
vp-kmod-3.0.1-4.18.0_372.32.1_1.el8.x86_64
# rpm -q ve_drv-kmod
ve_drv-kmod-3.0.1-4.18.0_372.32.1.el8_6.x86_64_1.el8.x86_64
```

If the version of the installed packages is newer than the version of the vp-kmod and ve_drv-kmod packages corresponding to the Linux kernel version you are using, you need to downgrade the vp-kmod and ve_drv-kmod packages.

3.3.2 Downgrade of the Latest vp-kmod and ve_drv-kmod Packages

Downgrade the vp-kmod and ve_drv-kmod packages by executing the following commands.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
# yum downgrade vp-kmod ve_drv-kmod
# /opt/nec/ve/sbin/start-all-ve-services
```

3.4 Confirm VE device path information file

If your model is defined on Table 3 (in 1.2.1 Hardware), please confirm the VE device path information file.

Confirm the VE device path information file using the following command:

```
# ls -l /etc/opt/nec/ve/veos/ve_nodes.json
-rw-r--r--. 1 root root 274  Feb 22  2023 /etc/opt/nec/ve/veos/ve_nodes.json
```

If the VE device path information file exists. Please go to section 3.5.

If VE device path information file does not exist, execute the following command:

```
# /opt/nec/ve/mmm/analysis/sbin/mmm-config-json.sh 2
```

After executing the above commands, reconfirm the VE device path information file and then execute the following restart ve-services commands:

```
# ls -l /etc/opt/nec/ve/veos/ve_nodes.json
-rw-r--r--. 1 root root 274  Feb 22  2023 /etc/opt/nec/ve/veos/ve_nodes.json
# /opt/nec/ve/sbin/terminate-all-ve-services
# /opt/nec/ve/sbin/start-all-ve-services
```

3.5 Status Check of the VEs

Confirm the status of the VEs is ONLINE with the vecmd command as the superuser. It can take a few minutes until it becomes ONLINE.

```
# /opt/nec/ve/bin/vecmd state get
Vector Engine MMM-Command v1.0.0
Command:
state -N 0 get
--------------------------------------------------------------------------------
VE0 [03:00.0] [ ONLINE ] Last Modif:2017/11/29 10:18:00
--------------------------------------------------------------------------------
Result: Success
```

⚠ Note

If "UNINITIALIZED" or "OFFLINE" is displayed, please wait for a short while.
If VE state does not become ONLINE after waiting, please restart services by executing the below commands.
 # /opt/nec/ve/sbin/terminate-all-ve-services
 # /opt/nec/ve/sbin/start-all-ve-services
And if "There is no executable ve card!" is displayed, please restart services in the same way.

If "UNAVAILABLE" is displayed, please restart the VHs by reboot command.
 # reboot

If "UNAVAILABLE" is still displayed after restarting, there may be a hardware failure. Please refer to the following guide and perform the separation.

✓ NEC Aurora Forum Documentation
"Vector Engine 2.0 Troubleshooting Guide"

3.6 Update of the VMC Firmware

Check whether update of the VMC firmware is required as follows:

```
# /opt/nec/ve/bin/vecmd fwup check
```

If the message "Updating VMCFW is required." is displayed, you need to update the VMC firmware. In this case, perform the following operations as the root user.

(1) Change the status of the VEs and update the firmware by vecmd commands

- When updating to VE1/VE2

```
# /opt/nec/ve/bin/vecmd state set off
# /opt/nec/ve/bin/vecmd state set mnt
# /opt/nec/ve/bin/vecmd fwup vmc aurora_MK10.bin
```

* The update of the firmware will take about a few minutes.

Reboot the VHs

```
# reboot
```

- When updating to VE3

```
# /opt/nec/ve/bin/vecmd state set off
# /opt/nec/ve/bin/vecmd fwup
```

* The update of the firmware will take about 10 minutes.

shutdown the VH and then turn it on.

Please note that this is different from VE1.

```
# shutdown -h now
```

(2) Log in as the root user

(3) Check the Status of the VEs

Confirm the status of the VEs is ONLINE. It can take a few minutes until it becomes ONLINE.

```
# /opt/nec/ve/bin/vecmd state get
Vector Engine MMM-Command v1.0.0
Command:
state -N 0 get
```

```
--------------------------------------------------------------------------------
VE0 [03:00.0] [ ONLINE ] Last Modif:2017/11/29 10:18:00
--------------------------------------------------------------------------------
Result: Success
```

⚠ Note

If "UNINITIALIZED" or "OFFLINE" is displayed, please wait for a short while.
And, if some VEs are not ONLINE or not appear, please shutdown VH and power on VH.

3.7 Command Compatible Packages Installation (Optional)

Starting with the March 2023 software release of SX-Aurora TSUBASA, commands ported from Linux will be prefixed with "ve-". If you need compability with commands ported from Linux before March 2023, please install the command compability package.

- VH for compiling and executing program

```
# yum install ve-command-name-compat-devel
```

- VH for executing program

```
# yum install ve-command-name-compat-runtime
```

- Frontend machine

```
# yum install ve-command-name-compat-front-end
```

3.8 Configuration of HugePages

VEOS, NEC MPI, and the ScaTeFS client use HugePages. Configure HugePages as below. If you have rebooted them at the update of the VMC firmware, this operation is not needed.

(1) Configure HugePages

```
# /opt/nec/ve/sbin/ve-set-hugepages
```

(2) Confirm HugePages setting

Confirm nr_hugepages nr_overcommit_hugepages are not 0.

```
# /opt/nec/ve/sbin/ve-set-hugepages -s
nr_hugepages:256
nr_overcommit_hugepages:23790
```

- Confirmation of the automatic configuration during booting

The systemd-sysctl service can configure HugePages during booting VH, but the configuration can be used exclusively this automatic configuration command (ve-set-hugepages) or systemd-sysctl. If the configurations of the HugePages or the overcommit HugePages are valid in /etc/sysctl.conf, /etc/sysctl.d etc., delete them. Please refer to section 4.12 for detail.

3.9 Start of the ScaTeFS Client (Optional)

If you have installed the ScaTeFS client, start it by rebooting the VHs. If you have rebooted them at the update of the VMC firmware, this operation is not needed.

```
# reboot
```

Now that the installation of the SX-Aurora TSUBASA software has completed, please proceed to Chapter4 to configure the software.

Chapter4 Software Configuration

This chapter explains how to configure the SX-Aurora TSUBASA software on the VHs before starting to use the SX-Aurora TSUBASA system.

The software configuration outline are shown in the following table.

Table 9 Software Configuration

Section	Requirements
4.1 Configuration of Operation Network	When using Operation Network
4.2 Specification of the License Server	When using SDK, NEC MPI (*), ScaTeFS, or NQSV (*) Version 2.13.0 or later is not required.
4.3 Configuration of the License Server	When using SDK, NEC MPI (*), ScaTeFS, or NQSV (*) Version 2.13.0 or later is not required
4.4 Configuration of InfiniBand HCA Relaxed Ordering (Optional)	When using model A412-8, B401-8, B403-8 or C401-8
4.5 Configuration of InfiniBand HCA PCIe credit number (Optional)	When using model B302-8, B401-8 or C401-8
4.6 Configuration of ScaTeFS	When using ScaTeFS
4.7 Configuration of NQSV	When using NQSV
4.8 Configuration of NEC MPI	When using NEC MPI
4.8.1 SELinux	When using NEC MPI and SELinux is enabled
4.8.2 Firewall	When using NEC MPI and Firewall is enabled
4.8.3 InfiniBand QoS	When using NEC MPI and IB QoS is enabled
4.8.4 NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)	When using NEC MPI and IB NVIDIA SHARP function

Section	Requirements
4.8.5 Configuration of persistent device name for HCAs	When using NEC MPI
4.8.6 Binding HCAs to VEs and Selecting Communication Method (Optional)	When using NEC MPI and model A412-8, B401-8, B403-8 or C401-8
4.8.7 Name Resolution of Hostname	When using NEC MPI
4.8.8 InfiniBand Adaptive Routing (Optional)	When using NEC MPI and enabling Adaptive Routing feature in the InfiniBand network
4.8.9 Partial Process Swapping (Optional)	When running NEC MPI program on VE from NQSV and using Partial Process Swapping
4.9 Setup of NEC Parallel Debugger	When using NEC Parallel Debugger in SDK
4.10 The Confirmation of the Virtual Memory Setting	When using ve_exec command
4.11 The Confirmation of the Memlock Resource Setting	When executing program
4.12 Configuration of HugePages	When using HugePages in VEOS (Accelerated I/O), NEC MPI, or the ScaTeFS client
4.13 Configuration for Partial Process Swapping	When using the function that save and restore the VE memory of suspended VE process
4.14 Configuration for Process accounting	When using the feature to get process accounting information on a VE process
4.15 How to Execute Programs on VEs	When executing program

4.1 Configuration of Operation Network

4.1.1 InfiniBand (IP over IB)

If you use InfiniBand (IP over IB) for the operation network, configure the ib0 network interface.

Please refer to Appendix C for the network configuration.

```
# nmcli connection add type infiniband autoconnect yes con-name ib0 ifname ib0
# nmcli connection mod ib0 ipv4.method manual ipv4.address "XX.XX.XX.XX/YY" (*)
```

(*) XX.XX.XX.XX is the IP address of IP over IB, and YY is the number of bits in the netmask.

In the case two or more InfiniBand HCA cards are installed, create the bond0 interface with the active-backup mode using two HCAs, and configure the bond0 network interface.

```
# nmcli connection add type bond con-name bond0 ifname bond0 mode active-backup
# nmcli connection mod bond0 ipv4.method disabled ipv6.method ignore
# nmcli connection add type infiniband autoconnect yes ifname ib0 master bond0
# nmcli connection add type infiniband autoconnect yes ifname ib1 master bond0
# nmcli connection mod bond0 +bond.options primary=ib0
# nmcli connection mod bond0 +bond.options miimon=100,updelay=100,downdelay=100
# nmcli connection mod bond0 ipv4.method manual ipv4.address "XX.XX.XX.XX/YY" (*)
```

(*) XX.XX.XX.XX is the IP address for IP over IB, and YY is the number of bits in the netmask.

4.1.2 Ethernet

If you use Ethernet for the operation network, configure the Ethernet interface for the operation network. Firstly, check the name of the Ethernet interface using the ifconfig or ip command, and then set it up.

```
# ifconfig -a
enp129s0f0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
       inet xx.xx.xx.xx  netmask yy.yy.yy.yy  broadcast zz.zz.zz.zz
       inet6 fe80::ec4:7aff:feea:d338  prefixlen 64  scopeid 0x20<link>
       ether 0c:c4:7a:ea:d3:38  txqueuelen 1000  (Ethernet)
       RX packets 215948  bytes 27415658 (26.1 MiB)
       RX errors 0  dropped 0  overruns 0  frame 0
       TX packets 943  bytes 475319 (464.1 KiB)
       TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

```
enp129s0f1: flags=4099<UP,BROADCAST,MULTICAST>  mtu 1500
       ether 0c:c4:7a:ea:d3:39  txqueuelen 1000  (Ethernet)
       RX packets 0  bytes 0 (0.0 B)
       RX errors 0  dropped 0  overruns 0  frame 0
       TX packets 0  bytes 0 (0.0 B)
       TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

# nmcli connection add type ethernet autoconnect yes con-name enp129s0f1 ifname enp129s0f1
# nmcli connection mod enp129s0f1 ipv4.method manual ipv4.address "XX.XX.XX.XX/YY" (*)
```

(*) enp129s0f1 is an example of the name of the Ethernet interface for the operation network. XX.XX.XX.XX is the IP address, and YY is the number of bits in the netmask.

4.1.3 Restart network

After the configuration, restart the network service.

In case RHEL/Rocky Linux 8.x:

```
# systemctl restart NetworkManager.service
# nmcli con up interface_name(*)
```

(*) Please specify the operation network interface name you've set.

4.2 Specification of the License Server

It is necessary to specify the license server on the VHs where the paid software is executed so that the software can obtain the license.

There are the following two methods to specify the license server.

- Specification with the Configuration File

  Specify the hostname and port number of the license server in the configuration file /opt/nec/aur_license/aur_license.conf as shown in Table 10.

Table 10 Parameters for Specifying the License Server

Configuration Item	Parameter Name	Value
Hostname of the License Server	License_server_host	Hostname
Receiving Port Number	License_server_port	Port Number

An example of the file aur_license.conf is as follows:

```
License_server_host=sv_host
```

```
License_server_port=7300
```

Please use this method for the ScaTeFS client.

- Specification with the Environment Variables

  It is possible to change the license servers for each software executed on VH (NEC MPI, C/C++ compiler and Fortran compiler) by specifying the license server with the environment variables as shown in Table 11.

Table 11 Environment Variables for Specifying the License Server

Configuration Item	Environment Variable
Hostname of the License Server	AURLIC_SERVER_HOSTNAME
Receiving Port number	AURLIC_SERVER_PORT

An example of the specification is as follows (bash):

```
$ export AURLIC_SERVER_HOSTNAME="sv_host"
$ export AURLIC_SERVER_PORT="7300"
```

The specification with the environment variables takes precedence over that with the configuration file.

4.3 Configuration of the License Server

If your SX-Aurora TSUBASA system is standalone, install and configure the license server on the VH. Otherwise, install and configure the license server on the management server. Please refer to "HPC Software License Management Guide" attached to the paid software for the setup of the license server.

4.4 Configuration of InfiniBand HCA Relaxed Ordering (Optional)

When you use model A412-8, B401-8, B403-8 or C401-8, you must enable PCIe Relaxed Ordering of InfiniBand HCA in order to increase in performance.

⚠ Note

- When you enable the Relaxed Ordering, please confirm whether middle wares and applications (software using InfiniBand Verbs API) you will run on SX-Aurora TSUBASA are compatible with Relaxed Ordering. Please contact to software venders whether the software can use with Relaxed Ordering configuration.

- NEC MPI (2.6.0 or later) and ScaTeFS are compatible with Relaxed Ordering.
- **Do not enable Relaxed Ordering unless model A412-8, B401-8, B403-8 or C401-8. Data corruption may occur.**

Please enable Relaxed Ordering for installed HCAs as below.

Please note that the XXXX shown in this procedure as "/dev/mst/mtXXXX_pciconf0" and "/dev/mtXXXX_pciconf1" depends on product model. Please set following value.

Model A412-8, B401-8 … 4123 (Ex. /dev/mst/mt4123_pciconf0)

Model B403-8, C401-8 … 4129 (Ex. /dev/mst/mt4129_pciconf0)

```
# mst start
# mlxconfig -d /dev/mst/mtXXXX_pciconf0 --yes set PCI_WR_ORDERING=1
…
Configurations:                          Next Boot       New
        PCI_WR_ORDERING                  per_mkey(0)     force_relax(1)

 Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.

# mlxconfig -d /dev/mst/mtXXX_pciconf1 --yes set PCI_WR_ORDERING=1
…
Configurations:                          Next Boot       New
        PCI_WR_ORDERING                  per_mkey(0)     force_relax(1)

 Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations
```

Please restart the system to take effect.

```
# reboot
```

Moreover, when "PCI_WR_ORDERING=0" is specified to mlxconfig command option, Relaxed Ordering will be disabled. When changed this setting, please restart the system as well.

After restarted the system, please confirm that the Relaxed Ordering is enabled as below.

```
# mst start
# mlxconfig -d /dev/mst/mtXXXX_pciconf0 query | grep PCI_WR_ORDER
        PCI_WR_ORDERING                  force_relax(1)
# mlxconfig -d /dev/mst/mtXXXX_pciconf1 query | grep PCI_WR_ORDER
        PCI_WR_ORDERING                  force_relax(1)
```

When PCI_WR_ORDERING is "force_relax(1)", it means Relaxed Ordering is enabled.

When PCI_WR_ORDERING is "per_mkey(0)", it means Relaxed Ordering is disabled.

4.5 Configuration of InfiniBand HCA PCIe credit number (Optional)

When you use model B302-8, B401-8 or C401-8, you must change InfiniBand HCA PCIe credit number from 16(default) to specified value in order to increase in performance.

The credit number depends on product model. Please set following value.

Model B302-8 … 32

Model B401-8 … 32

Model C401-8 … 36

YY means credit number.

⚠ Note

- Configiration of HCA PCIe credit number is available when HCA Firmware version is 20.28.2006 or later.
- Do not set PCIe clredit number other than above values.

Please change the credit number for installed HCAs as below.

Please note that the XXXX shown in this procedure as "/dev/mst/mtXXXX_pciconf0" and "/dev/mtXXXX_pciconf1" depends on product model. Please set following value.

Model B302-8, B401-8 … 4123 (Ex. /dev/mst/mt4123_pciconf0)

Model C401-8 … 4129 (Ex. /dev/mst/mt4129_pciconf0)

```
# mst start
# mlxconfig -d /dev/mst/mtXXXX_pciconf0 --yes set ADVANCED_PCI_SETTINGS=1
…
Configurations:                          Next Boot      New
        ADVANCED_PCI_SETTINGS             False(0)       True(1)

 Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.

# mlxconfig -d /dev/mst/mtXXXX_pciconf0 --yes set MAX_ACC_OUT_READ=YY
…
Configurations:                          Next Boot      New
        MAX_ACC_OUT_READ                  0              YY

 Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.
```

```
# mlxconfig -d /dev/mst/mtXXXX_pciconf1 --yes set ADVANCED_PCI_SETTINGS=1
…
Configurations:                          Next Boot      New
        ADVANCED_PCI_SETTINGS              False(0)      True(1)

 Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.

# mlxconfig -d /dev/mst/mtXXX_pciconf1 --yes set MAX_ACC_OUT_READ=YY
…
Configurations:                          Next Boot      New
        MAX_ACC_OUT_READ                   0             YY

 Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.
```

Please restart the system to take effect.

```
# reboot
```

After reboot, please check if the PCIe credit number is changed correctly as below.

```
# mst start
# mlxconfig -d /dev/mst/mtXXXX_pciconf0 query | grep ADVANCED_PCI_SETTINGS
        ADVANCED_PCI_SETTINGS              True(1)
# mlxconfig -d /dev/mst/mtXXXX_pciconf0 query | grep MAX_ACC_OUT_READ
        MAX_ACC_OUT_READ                   YY
# mlxconfig -d /dev/mst/mtXXXX_pciconf1 query | grep ADVANCED_PCI_SETTINGS
        ADVANCED_PCI_SETTINGS              True(1)
# mlxconfig -d /dev/mst/mt4123_pciconf1 query | grep MAX_ACC_OUT_READ
        MAX_ACC_OUT_READ                   YY
```

The ADVANCED_PCI_SETTINGS should be "True(1)" and MAX_ACC_OUT_READ should be specified value for each HCA.

4.6 Configuration of ScaTeFS

When you use the ScaTeFS client, configure the ScaTeFS client on the VHs referring to "Chapter 6 Setting the Linux client" of "NEC Scalable Technology File System (ScaTeFS) Administrator's Guide".

The ScaTeFS client uses HugePages, refer to section 4.12 for the configuration.

As for setup of the ScaTeFS server, please refer to "NEC Scalable Technology File System (ScaTeFS) Administrator's Guide".

4.7 Configuration of NQSV

When you use NQSV, configure the job server and client referring to "NEC Network Queuing System V (NQSV) User's Guide [Management]".
Also, create and configure the file /etc/opt/nec/nqsv/resource.def to define device resources on each job server host (VH) according to "5.4 HCA Assignment Feature" of "NEC Network Queuing System V (NQSV) User's Guide [JobManipulator]".

As for setup of the NQSV server (the batch server, scheduler, and accounting functions), please refer to "NEC Network Queuing System V (NQSV) User's Guide [Introduction]".

4.8 Configuration of NEC MPI

This section describes the settings needed on all VHs where MPI programs are launched.
The NEC MPI uses HugePages, refer to section 4.12 for the configuration.

4.8.1 SELinux

If SELinux is enabled, please perform the following setting. Otherwise, it is not needed.
Please set the boolean use_nfs_home_dirs to 1 if MPI programs are interactively executed over multiple VHs and the home directory is mounted via NFS or ScaTeFS.

```
# setsebool -P use_nfs_home_dirs 1
```

The message "use_nfs_home_dirs --> on" is displayed with the following command if the setting is successful.

```
# getsebool use_nfs_home_dirs
use_nfs_home_dirs --> on
```

4.8.2 Firewall

If the firewall is enabled, please perform the following setting. Otherwise, it is not needed.
NEC MPI uses TCP/IP ports from 25257 to 25266 by default in order to accept connections. Therefore, please execute the following commands to open the ports.

```
# firewall-cmd --zone=public --permanent --add-port=25257-25266/tcp
# firewall-cmd --reload
```

The following command displays the open ports.

```
# firewall-cmd --list-port --zone=public
25257-25266/tcp
```

TCP/IP ports used by NEC MPI can be changed with the environment variable NMPI_PORT_RANGE. The following example specifies that NEC MPI uses the TCP/IP ports range from 25257 to 25266.

```
# export NMPI_PORT_RANGE=25257:25266
```

If you run urgent jobs by using the Partial Process Swapping feature (cf. section 4.13) under NQSV, additional ports can be necessary. When the number of jobs on one VH (including jobs suspended by the urgent jobs) exceeds 10, please open at least as many number of ports as the total number of the jobs by firewall-cmd command and set the port numbers to the environment variable NMPI_PORT_RANGE.

4.8.3 InfiniBand QoS

If Quality of Service (QoS) is enabled for the InfiniBand Subnet Manager, please set the service level used by NEC MPI. Otherwise, this setting is not needed.

Set the parameter ib_qos_venode to the value of the service level in the configuration file /etc/opt/nec/ve/mpi/necmpi.conf as follows. If there is no setting, the service level 0 is applied, which is the highest QoS level. Note that lines beginning with "#" are treated as comments and ignored.

```
# Ex) In the case that the service level 2 is set
ib_qos_venode   2
```

4.8.4 NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)

NEC MPI supports NVIDIA SHARP. This function requires configuration of the Aggregation Manager. The Aggregation Manager runs on the same host as the InfiniBand Subnet Manager. The versions of the Aggregation Manager and SHARP library have to be identical. The SX-Aurora TSUBASA supports only following versions.

MLNX_OFED version	SHARP version
23.10-3.2.2.0	3.5.2

⚠ Note

If you use NVIDIA UFM as the Subnet Manager, the version of SHARP provided by NVIDIA UFM may not match the version of SHARP provided by MLNX_OFED. Please contact NEC support for compatibility.

- Configuration of the VH

  Configuration of firewall is required on each VHs. Please configure the firewall to accept communications from Subnet Manager. The following is an example.

```
# firewall-cmd --add-source=<IPoIB address of Subnet Manager> --zone=trusted --permanent
# firewall-cmd --add-source=<IPoIB address of Subnet Manager> --zone=trusted
```

- Configuration of the Aggregation Manager

  As for the configuration of the Aggregation Manager, please refer to documents on the following NVIDIA official home page:

  https://docs.nvidia.com/networking/category/mlnxsharp

  If the firewall is running on the host where the Aggregation Manager runs, open TCP port 6126 on the host as follows, because the Aggregation Manager communicates with the SHARP daemons via the TCP connection.

```
# firewall-cmd --add-port=6126/tcp --permanent
# firewall-cmd --add-port=6126/tcp
```

  If the host where the Aggregation Manager runs has multiple network interfaces any of which cannot directly communicate with the VHs, it is necessary to explicitly specify the network interface the Aggregation Manager uses. Set the parameter smx_sock_interface to the name of the network interface in the configuration file /opt/mellanox/sharp/conf/sharp_am.cfg on the host. The following example specifies ib0 as the network interface.

```
smx_sock_interface ib0
```

4.8.5 Configuration of persistent device name for HCAs

The default device name of HCAs (mlx5_X) may change when updating MLNX_OFED or updating the VH's BIOS. If you use InfiniBand with the SX-Aurora TSUBASA, please create /etc/udev/rules.d/60-rdma-persistent-naming.rules with the following content for configuring persistent device name.

```
ACTION=="add", SUBSYSTEM=="infiniband", PROGRAM="rdma_rename %k NAME_PCI"
```

Please reboot the VH after configuration.

```
# reboot
```

⚠ Note

By following this procedure, the HCA device name will be changed to ibpXXs0 (where XX is the PCI bus number of the HCA). If you are using software other than SX-Aurora TSUBASA software that directly specifies the HCA device name, you may need to modify the settings

4.8.6 Binding HCAs to VEs and Selecting Communication Method (Optional)

When you use the model A412-8, B401-8, B403-8 or C401-8, please add the following line to the configuration file /etc/opt/nec/ve/mpi/necmpi.conf, which may help improve application performance. It is unnecessary for the other models.

- for A412-8 B401-8

```
ib_adapter_name_list_pve  0=ibp69s0:1  1=ibp69s0:1  2=ibp69s0:1  3=ibp69s0:1  4=ibp1s0:1
5=ibp1s0:1 6=ibp1s0:1 7=ibp1s0:1
```

- for B403-8

```
ib_adapter_name_list_pve  0=ibp78s0:1  1=ibp78s0:1  2=ibp78s0:1  3=ibp78s0:1  4=ibp1s0:1
5=ibp1s0:1 6=ibp1s0:1 7=ibp1s0:1
```

- for C401-8

```
ib_adapter_name_list_pve  0=ibp78s0:1  1=ibp78s0:1  2=ibp78s0:1  3=ibp78s0:1  4=ibp1s0:1
5=ibp1s0:1 6=ibp1s0:1 7=ibp1s0:1
dma_vh_memcpy     off
ib_vh_memcpy_send  off
ib_vh_memcpy_recv  off
```

The line starting with "ib_adapter_name_list_pve" defines HCAs to be used for MPI communication for each VE. After keyword "ib_adapter_name_list_pve", please place the list of HCAs used by each VE in the syntax of "<VE0>=<HCA1>[,<HCA2> ...] [<VE1>=<HCA1>[,<HCA2> ...] ...]" in a single line. Each line staring with "dma_vh_memcpy", "ib_vh_memcpy_send" or "ib_vh_memcpy_recv" defines the usage of VH memory copy at MPI communication. Please set all these to off in the case of C408-1. Note that lines beginning with "#" are treated as comments and ignored.

4.8.7 Name Resolution of Hostname

The host used for MPI execution must be able to resolve the hostname to an IP address on the host. This setting is also required for single node execution. Please configure on the host that the hostname can be resolved to an IP address by using /etc/hosts, DNS and so on.

4.8.8 InfiniBand Adaptive Routing (Optional)

If Adaptive Routing feature is enabled in the InfiniBand network, please add the following line to the configuration file /etc/opt/nec/ve/mpi/necmpi.conf.

```
# Adaptive Routing enabled in InfiniBand network
ib_ar 1
```

The line with the keyword "ib_ar" followed by a blank and "1" tells NEC MPI that Adaptive Routing feature is enabled, and NEC MPI will select the execution for it. If there is no such setting, NEC MPI assumes Adaptive Routing is not enabled. Note that lines beginning with "#" are treated as comments and ignored. And if the setting for NEC MPI is conflicting with the actual InfiniBand Adaptive Routing setting, there is a possibility of problems such as MPI communication performance degradation.

4.8.9 Partial Process Swapping (Optional)

In the execution through NQSV, if normal jobs should be suspended with the Partial Process Swapping feature and Non Swappable Memory of MPI processes running on VEs should be released, please add the following line to the configuration file /etc/opt/nec/ve/mpi/necmpi.conf. Also, if off is specified for "ib_vh_memcpy_send" and "ib_vh_memcpy_recv" in 4.8.6, disable them by adding # to the beginning of

these two lines.

```
# Release of Non Swappable Memory
swap_on_hold 1
```

The line with the keyword "swap_on_hold" followed by a blank and "1" tells NEC MPI that Non Swappable Memory should be released, when suspending. Note that lines beginning with "#" are treated as comments and ignored. Please note that it is required to use Hydra as the process manager for NQSV batch queues used to run MPI programs. If MPD as the process manager is used, there is a possibility that MPI programs cannot be executed in an urgent request started after a preceding normal request suspended.

4.9 Setup of NEC Parallel Debugger

NEC Parallel Debugger is an Eclipse PTP plugin for debugging. This section describes the settings needed on the hosts where the Eclipse PTP is launched and on the VHs where programs debugged by NEC Parallel Debugger are launched.

4.9.1 Installation of Eclipse PTP

Download the latest version of the Eclipse PTP distribution, Oxygen version 3, from the Eclipse official site onto the host where Eclipse PTP is launched.

Then expand the downloaded file to an arbitrary install directory (/INST-DIR in the following example).

```
# tar zxvf eclipse-parallel-oxygen-3-linux-gtk-x86_64.tar.gz -C /INST-DIR
```

The executable of Eclipse PTP is placed as /INST-DIR/eclipse/eclipse. Please add the path to the environment variable PATH.

4.9.2 Installation of the NEC Parallel Debugger Plugin

Put the NEC Parallel Debugger plugin onto the directory "plugins" under the install directory.

```
# cp /opt/nec/ve/npdb/plugins/com.nec.ParallelDebugger_*.jar /INST-DIR/eclipse/plugins/
```

4.9.3 Installation of the Necessary Software for Eclipse PTP

The following software has to be installed for Eclipse PTP to work. Please download and install the latest version onto the host where Eclipse PTP is launched from a download site such as the CPAN Search Site.

- Test-Simple
- Devel-GDB
- Expect
- IO-Tty

- Test-Simple

  Expand the downloaded file under an arbitrary directory, and make and install it.

```
# tar zxvf Test-Simple-1.302183.tar.gz
# cd Test-Simple-1.302183
# perl Makefile.PL
# make
# make test
# make install
```

- Devel-GDB

  Expand the downloaded file under an arbitrary directory, and make and install it.

```
# tar zxvf Devel-GDB-2.02.tar.gz
# cd Devel-GDB-2.02
# perl Makefile.PL
# make
# make test
# make install
```

- IO-Tty

  Expand the downloaded file under an arbitrary directory, and make and install it.

```
# tar zxvf IO-Tty-1.16.tar.gz
# cd IO-Tty-1.16
# perl Makefile.PL
# make
# make test
# make install
```

- Expect

  Expand the downloaded file under an arbitrary directory, and make and install it.

```
# tar zxvf Expect-1.35.tar.gz
# cd Expect-1.35
# perl Makefile.PL
# make
# make test
# make install
```

4.9.4 Configuration of the Firewall

If the firewall is enabled on the VHs, please perform the following setting. Otherwise, it is not needed.

Scalable Debug Manager (SDM), which is a component of Eclipse PTP working with NEC Parallel Debugger, uses the TCP/IP ports from 50000 through 50079 by default in order to accept connections. Therefore, please execute the following commands to open the ports.

```
# firewall-cmd --zone=public --permanent --add-port=50000-50079/tcp
# firewall-cmd --reload
```

The following command displays the open ports.

```
# firewall-cmd --list-port --zone=public
50000-50079/tcp
```

TCP/IP ports used by SDM can be changed with the environment variable NPDB_SDM_PORTRANGE. The following example specifies that SDM uses the TCP/IP ports range from 51000 through 51100.

```
# export NPDB_SDM_PORTRANGE=51000:51100
```

4.10 The Confirmation of the Virtual Memory Setting

The ve_exec command, which is a part of VEOS, uses over 1TB of the VH virtual address space using the overcommit function. VEOS requires that the value of the parameter vm.overcommit_memory be other than 2 to enable the overcommit function. Please confirm the value of vm.overcommit_memory as follows.

```
# sysctl vm.overcommit_memory
vm.overcommit_memory = 0
```

⚠ Note

No VH physical memory is allocated to the VH virtual address space larger than 1TB.

4.11 The Confirmation of the Memlock Resource Setting

The memlock resource needs to be set to "unlimited". This setting is applied automatically when you install VEOS. Confirm it with the following command.

```
# ulimit -l
unlimited
```

Changed memlock setting is not reflected immediately. If the printed value is not "unlimited", please log out and log in back the VH, then re-confirm it.

4.12 Configuration of HugePages

VEOS (Accelerated I/O, Partial Process Swapping), NEC MPI, and the ScaTeFS client use HugePages. The automatic HugePages configuration command sets the number of HugePages and the number of overcommit HugePages (how large HugePages can be usable if it is available) automatically. The number of HugePages depends on the environment of VH, and the number of overcommit HugePages is the half of VH's total memory by default.

- Please refer to section 4.13 regarding Partial Process Swapping.
- If NQSV socket scheduling is required, refer to the NEC Network Queuing System V (NQSV) Users Guide [Administration], Chapter.18 Socket Scheduling, 18.1.3 Memory allocation policy. When you use the socket scheduling with membind policy, the number of HugePages should be "the number of VH numa nodes" times larger then the one without socket scheduling. Validate MEMBIND option in /etc/opt/nec/ve/veos/ve-hugepages.conf.

```
# vi /etc/opt/nec/ve/veos/ve-hugepages.conf
MEMBIND=YES
```

- When non-SX-Aurora TSUBASA software requires HugePages, set the number of HugePages in /etc/opt/nec/ve/veos/ve-hugepages.conf.

```
# vi /etc/opt/nec/ve/veos/ve-hugepages.conf
ADDITIONAL_HUGE_PAGES=<the number of Hugepages>
```

Please refer to the VEOS document "How to execute VE program" for detail.

Execute the configuration command again in order to update the values when you change the options of the automatic HugePages configuration command or packages for NEC MPI and ScaTeFS Clinent.

```
# /opt/nec/ve/sbin/ve-set-hugepages
```

The configuration of HugePages is set during the bootstrap automatically. Please note that you don't specify the options for HugePages and overcommit HugePages in /etc/sysctl.conf, /etc/sysctl.d, etc. because systemd-sysctl and this command work exclusively regarding HugePages.

> **⚠ Note**
>
> If HugePages (nr_hugepages) option and overcommit Hugepages (nr_overcommit_hugepages) option exist in the configuration file of systemd-sysctl.service (sysctl commtnd), the values of sysctl.service may be set. Please refer to the Linux manpage sysctl.d(5) regarding configuration files of systemd-sysctl.service.

4.13 Configuration for Partial Process Swapping

Partial Process Swapping is the feature to swap out VE memory of processes to VH memory or files and free VE memory. The VE memory freed by the swapping out is used for running other processes. After the run of the processes completes, execution of the process being suspended can be resumed by swapping in the VH memory previously swapped out to VE memory.

If you do not use Partial Process Swapping, skip this section.

4.13.1 Configuration of veswap.option

If you want to set the same values for all VEs, edit /etc/opt/nec/ve/veos/ve-os-launcher.d/veswap.options. If you want to set the different values for each VE, make /etc/opt/nec/ve/veos/ve-os-launcher.d/N directory where N is VE node number. Then, create veswap.options under the directory.

The contents of veswap.options are below. If you specify ve-os-launcher@*=--ve-swap-mem-max=, VE memory is swap out VH memory. If you specify both ve-os-launcher@*=--ve-swap-file-path= and ve-os-launcher@*=--ve-swap-file-max=, VE memory is swapped out to a file.

Table 12 The configuration in veswap.option file

Option	Value	Description	Note
ve-os-launcher@*=--ve-swap-mem-max=	Integer (MB)	The maximum size of VE memory to swap out to VH memory	The value is for one VE.
ve-os-launcher@*=--ve-swap-file-path=	String	The path of a file to which VE memory is swap out	A host name and a VE node number are added to the specified path.
ve-os-launcher@*=--ve-swap-file-max=	Integer (MB)	The maximum size of VE memory to swap out to a file	The value is for one VE.
ve-os-launcher@*=--ve-swap-file-user=	String	The owner of the swap file	The default value is "root", if the option is not specified.
ve-os-launcher@*=--ve-swap-file-group=	String	The group of the swap file	The default value is "root", if the option is not specified.

Provide examples with the following four case.

1. Swap out to VH memory
2. Swap out to a local file system
3. Swap out to a shared file system

4. Swap out to VH memory, and file systems on two NVMe drives parallelly

Example 1: Swap out to VH memory

Add the below line to veswap.options file.

```
ve-os-launcher@*=--ve-swap-file-max=49152
```

In this example, maximally 48GB (49,152MB) VE memory can be swapped out to VH memory for each VE.

Example 2: Swap out to a local file system

Add the below line to veswap.options file.

```
ve-os-launcher@*=--ve-swap-file-path=/var/opt/nec/ve/veos/swap-file
ve-os-launcher@*=--ve-swap-file-max=49152
```

In this example, maximally 48GB (49,152MB) VE memory can be swapped out to the file whose path is /var/opt/nec/ve/veos/swap-file_<hostname>.N where <hostname> is a host name and N is a VE node number.

Example 3: Swap out to a local file system

This is an example where the shared file system is mounted on /mnt/share, and the owner and group of the /mnt/share/veswap directory are the "veswap" user and the "veswap" group.

Add the below line to veswap.options file.

```
ve-os-launcher@*=--ve-swap-file-path=/mnt/share/veswap/swap-file
ve-os-launcher@*=--ve-swap-file-max=49152
ve-os-launcher@*=--ve-swap-file-user=veswap
ve-os-launcher@*=--ve-swap-file-group=veswap
```

In this example, maximally 48GB (49,152MB) VE memory can be swapped out to the file whose path is /mnt/share/veswap/swap-file_<hostname>.N where <hostname> is the host name and N is a VE node number. VEOS will create the file with privilege

of "veswap" user and "veswap" group.

Example 4: Swap out to VH memory, and file systems on two NVMe drives parallelly
In this example, maximally 48GB (49,152MB) VE memory can be swapped out to VH memory for each VE, and maximally 48GB (49,152MB) VE memory can be swapped out to the files created in file systems on NVMe drives parallelly.
Assuming a VH with 8 VEs and 2 NVMe drives, from VE0 to VE3 are swapped out to a file system on one NVMe drive, and from VE4 to VE7 are swapped out to a file system on another NVMe drive.
A file system is created on each of the two NVMe drives and mounted on /mnt/nvme0 and /mnt/nvme1.
Please check /etc/opt/nec/ve/veos/ve-os-launcher.d/veswap.options and make sure all settings are commented out.

Please create the 8 directories to store configuration file by executing the below command.

```
# for d in `seq 0 7`; do mkdir /etc/opt/nec/ve/veos/ve-os-launcher.d/$d/; done
```

Please create the following 4 files.

```
/etc/opt/nec/ve/veos/ve-os-launcher.d/0/veswap.options
/etc/opt/nec/ve/veos/ve-os-launcher.d/1/veswap.options
/etc/opt/nec/ve/veos/ve-os-launcher.d/2/veswap.options
/etc/opt/nec/ve/veos/ve-os-launcher.d/3/veswap.options
```

The contents are as follows.

```
ve-os-launcher@*=--ve-swap-mem-max=49152
ve-os-launcher@*=--ve-swap-file-path=/mnt/nvme0/swap-file
ve-os-launcher@*=--ve-swap-file-max=49152
```

Please create the following 4 files.

```
/etc/opt/nec/ve/veos/ve-os-launcher.d/4/veswap.options
/etc/opt/nec/ve/veos/ve-os-launcher.d/5/veswap.options
/etc/opt/nec/ve/veos/ve-os-launcher.d/6/veswap.options
/etc/opt/nec/ve/veos/ve-os-launcher.d/7/veswap.options
```

The contents are as follows.

```
ve-os-launcher@*=--ve-swap-mem-max=49152
```

```
ve-os-launcher@*=--ve-swap-file-path=/mnt/nvme1/swap-file
ve-os-launcher@*=--ve-swap-file-max=49152
```

With this setting, the paths of the swap-out destination files will be /mnt/nvme0/swap-file_<hostname>.N and /mnt/nvme1/swap-file_<hostname>.N, <hostname> is the hostname and N is the VE node number.

4.13.2 Restarting VEOS

Execute the automatic HugePages configuration command and restart VEOS as follows in order to load the configuration.

```
# /opt/nec/ve/sbin/ve-set-hugepages
# systemctl restart ve-os-launcher@*
```

Please confirm Partial Process Swapping is enabled using the veswap command with the -s option. If Partial Process Swapping is enabled, the values of the "SWAPPED" column are all "0". If Partial Process Swapping is disabled, the values of "SWAPPED" columns are "N/A". If the veswap command shows "Failed to get online VE Nodes: No such file or directory", VEOS is restarting. In this case, wait a few seconds, and retry it.

```
# /opt/nec/ve/bin/veswap -s
    VE     SWAPPED
    VE7          0
    VE6          0
    VE5          0
    VE4          0
    VE2          0
    VE3          0
    VE1          0
    VE0          0
```

NQSV can execute urgent jobs by using the Partial Process Swapping feature. For details on the setting method, refer to "5.6 Suspend Jobs Using VEs" of "NEC Network Queuing System V (NQSV) User's Guide [JobManipulator]".

4.14 Configuration for Process accounting

Process accounting is the feature to get process accounting information on a VE process. When process accounting is enabled, the process accounting information is written to process accounting file as each process on it terminates.

4.14.1 Start Process accounting service

If you use Process accounting, start psacct-ve service with executing this command.

```
# for i in `seq 0 7`; do if [ -e /dev/veslot$i ]; then  systemctl enable psacct-ve@$i;
systemctl restart psacct-ve@$i; fi done
```

If psacct-ve service is active, process accounting information is recorded when VE process terminates. You can read this information with ve-lastcomm command. Please choice a process accounting file which corresponds to the node number.

```
# /opt/nec/ve/bin/ve-lastcomm -f /var/opt/nec/ve/account/pacct_*
```

For example, this is how to read node number 0 process accounting file.

```
# /opt/nec/ve/bin/ve-lastcomm -f /var/opt/nec/ve/account/pacct_0
```

4.14.2 Stop Process accounting service

If you don't use Process accounting, stop psacct-ve service with executing this command.

```
 # for i in `seq 0 7`; do systemctl stop psacct-ve@$i; systemctl disable psacct-ve@$i $i;
done
```

4.15 How to Execute Programs on VEs

There are two ways to execute programs on VEs as below.

- Execution with the ve_exec command
- Execution without the ve_exec command

- Execution with the ve_exec Command

  It is necessary to use the ve_exec command to execute programs on VEs in the SX-Aurora TSUBASA system. To omit the ve_exec command in command lines, please refer to the following explanation.

```
[Example]
In the case of an MPI program:
  $ mpirun -host 0 -ve 0-1 -np 16 /opt/nec/ve/bin/ve_exec ./mpi.lm
     mpi.lm is the file name of the program to execute.
```

```
In the case of a non-MPI program:
  $ /opt/nec/ve/bin/ve_exec ./nonmpi.lm
    nonmpi.lm is the file name of the program to execute.
```

- Execution without the ve_exec Command

It is possible to execute programs on VEs without specifying the ve_exec command in command lines, by configuring the Linux binfmt_misc capability, which is the default software development environment on VEOS 1.0.3 or later. This is commonly available for batch jobs, interactive jobs, and interactive executions.

```
In the case of an MPI program:
  $ mpirun -host 0 -ve 0-1 -np 16 ./mpi.lm
    mpi.lm is the file name of the program to execute.


In the case of a non-MPI program:
  $ ./nonmpi.lm
    nonmpi.lm is the file name of the program to execute.
```

Chapter5　Update

This chapter explains how to update the SX-Aurora TSUBASA software installed on the VHs.
When you have the local yum repository, you can update the repository.

Update of the paid software requires the PP support contract.
It is available to update the free software for all users.

You can update the software installed on VHs with following procedure.

⚠ Note

When updating OS, MLNX_OFED may need to be updated from version 4.x to 23.x (see Section 2.3).

In this case, some features are not compatible between MLNX_OFED versions 4.x and 23.x, so an executable file that statically links NEC MPI, it is a default setting, in an environment before update cannot be executed in an environment after update.

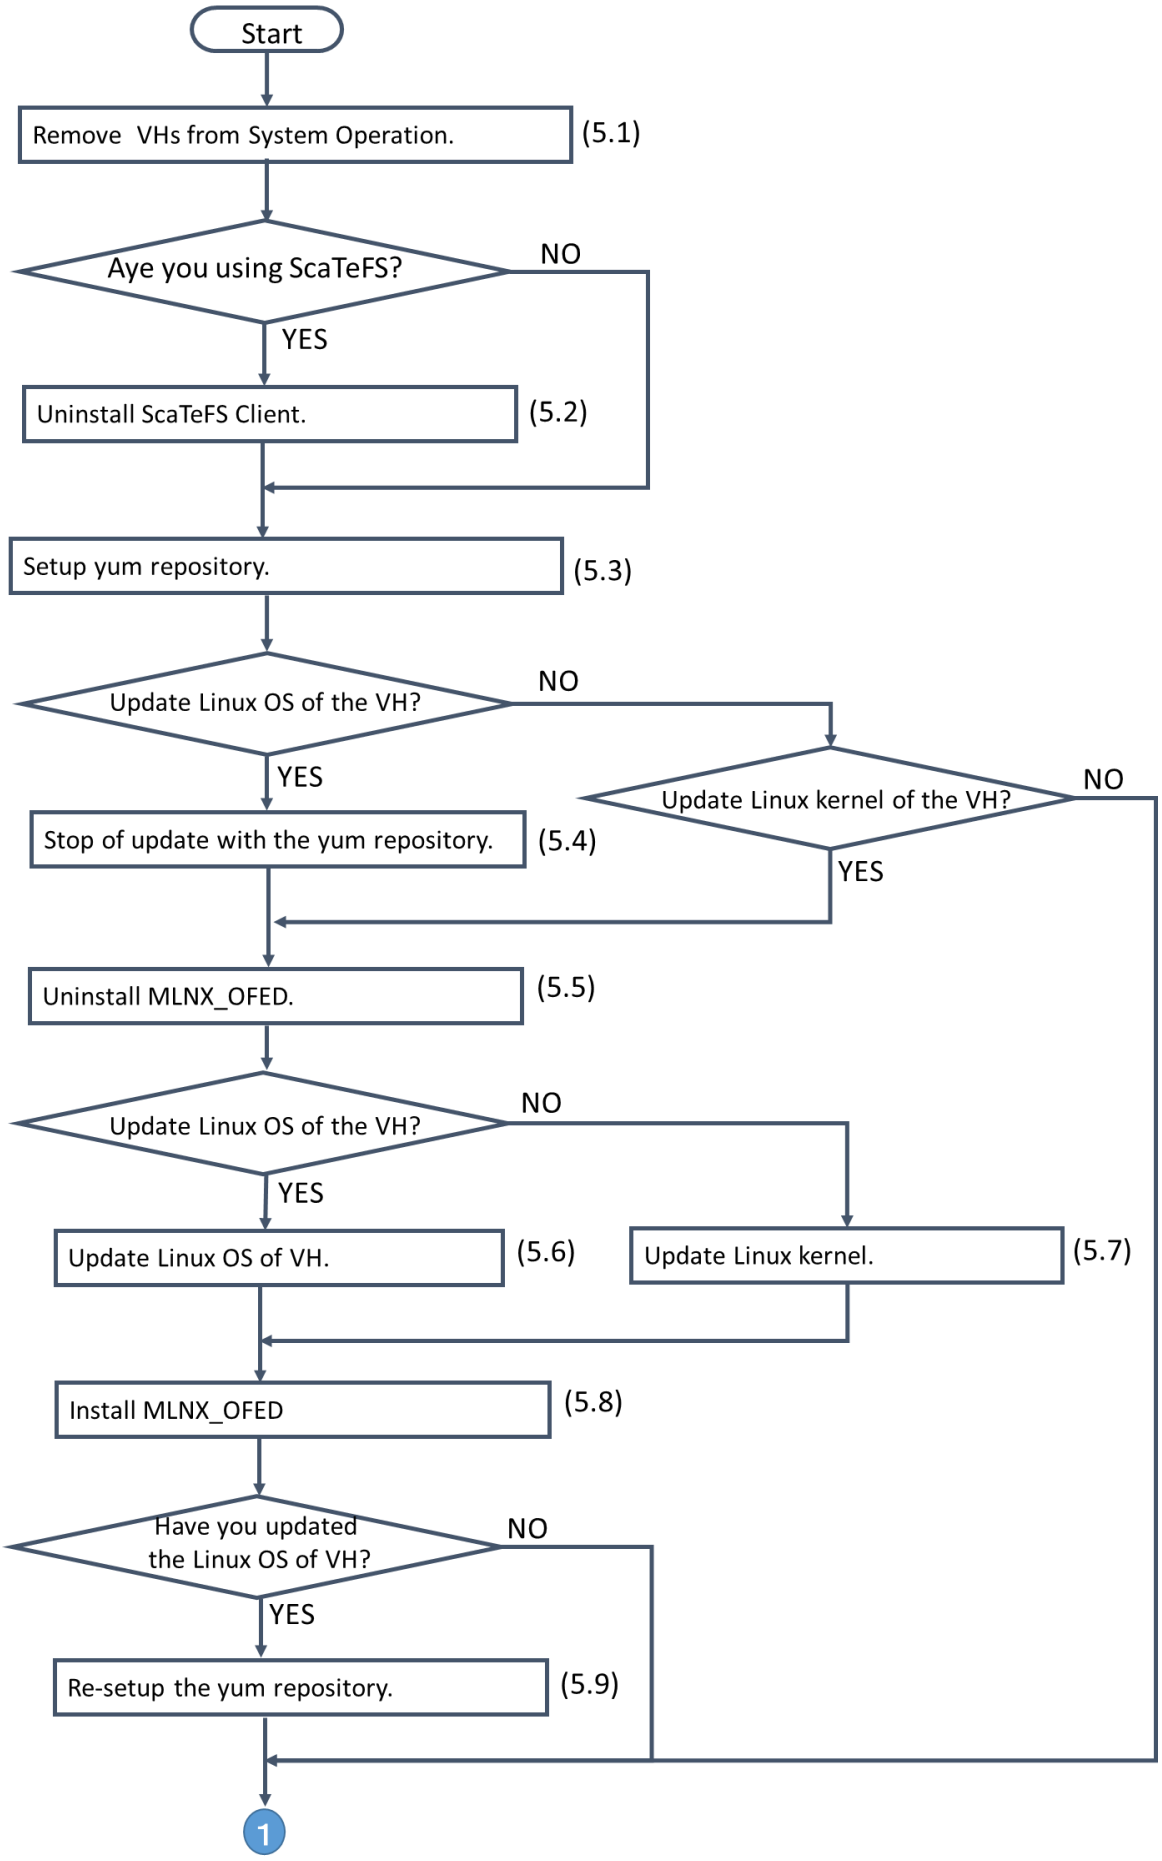
Start
Remove VHs from System Operation.
(5.1)
Aye you using ScaTeFS?
NO
YES
Uninstall ScaTeFS Client.
(5.2)
Setup yum repository.
(5.3)
Update Linux OS of the VH?
NO
YES
Update Linux kernel of the VH?
NO
YES
Stop of update with the yum repository.
(5.4)
Uninstall MLNX_OFED.
(5.5)
Update Linux OS of the VH?
NO
YES
Update Linux OS of VH.
(5.6)
Update Linux kernel.
(5.7)
Install MLNX_OFED
(5.8)
Have you updated the Linux OS of VH?
NO
YES
Re-setup the yum repository.
(5.9)
1

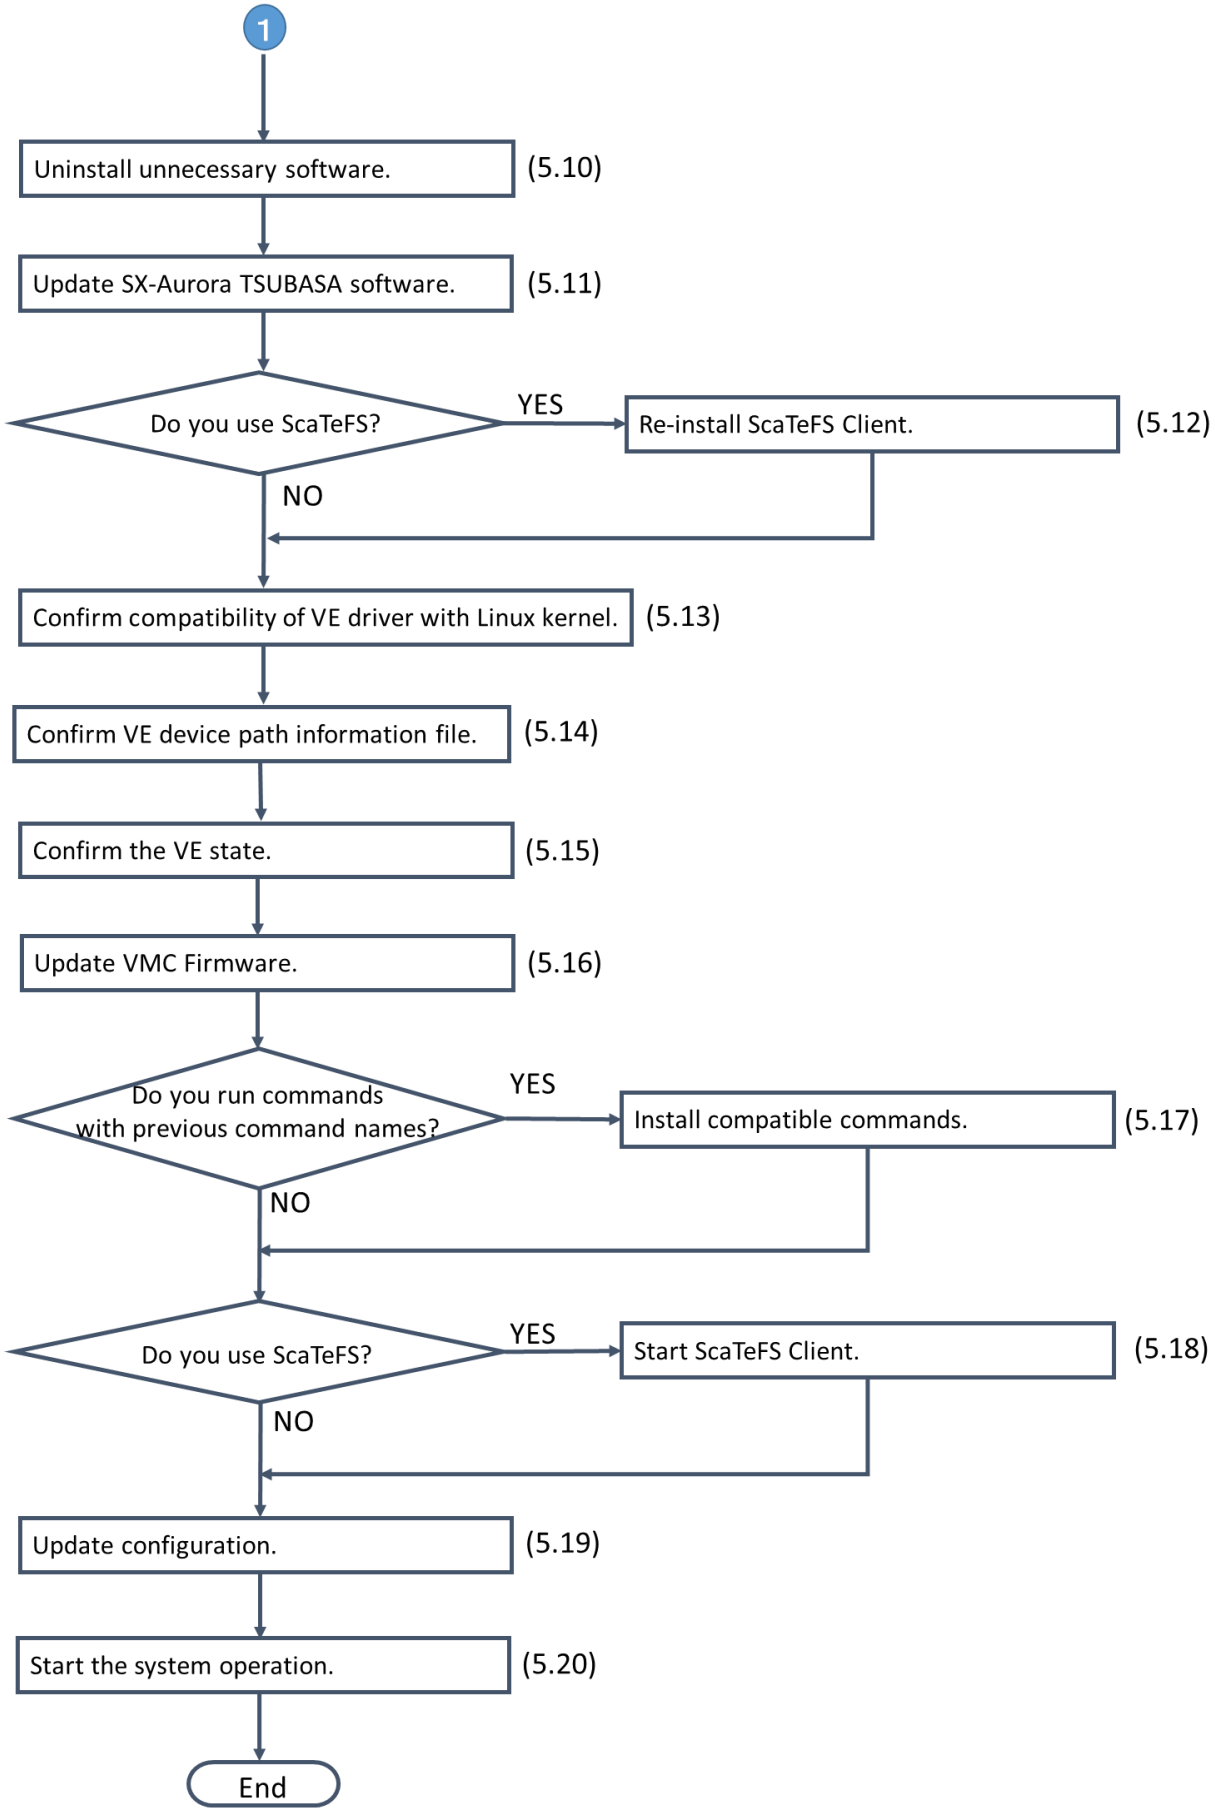

Figure 9 Software Update

5.1 Removal of VHs from System Operation

5.1.1 Disconnection of VHs from the Job Scheduler

If you use the job scheduler, disconnect the VHs targeted for update from the job scheduler.

1. Unbind the job server on the VHs from the queue which is being operated using the unbind execution_queue subcommand of the qmgr(1M) command on the host where NQSV/Client is installed. For example, the job server with job server ID 100 is unbound from the bq queue by the following operation.

```
$ qmgr -Pm
Mgr: unbind execution_queue job_server bq job_server_id = 100
```

   Use the unbind interactive_queue subcommand for interactive queues.

2. Make sure that the qstat(1) command with the -J option does not display the job server ID unbound in Step 1 in the JSVNO field on the host where NQSV/Client is installed, to confirm that no jobs exist on the job server.

```
$ qstat -J -Pm
```

   If any jobs exist on the job server, perform one of the following operations according to your operation policy.

   - Wait for the jobs to complete execution.
   - Re-run the jobs on different hosts with the qrerun command.
   - Delete the jobs with the qdel command.

3. Stop the job server and launcher on the VHs using the systemctl command.

```
# systemctl stop nqs-jsv.target
```

5.1.2 Removal of the VHs from Target of Monitoring

If you use the monitoring software (Zabbix or Ganglia+Nagios) for the VHs, please place the VHs into the maintenance mode.

5.2 Uninstallation of the ScaTeFS Client (Optional)

If you use the ScaTeFS client on the VHs, stop the ScaTeFS client service and uninstall the ScaTeFS client according to the following procedure.

1. If a ScaTeFS file system on the VHs targeted for update is exported via NFS, unmount the ScaTeFS file system on all the NFS clients, then go to Step 2.
   Otherwise, go to Step 3.

2. Stop the nfs-server service on the VHs.

```
# systemctl stop nfs-server
```

3. If a ScaTeFS file system on the VHs is exported as a share via Samba, stop accessing the Samba share on all the CIFS clients such as Windows, then go to Step 4.
   Otherwise, go to Step 5.

4. Stop the smb service and nmb service on the VHs.

```
# systemctl stop smb
# systemctl stop nmb
```

5. Unmount the ScaTeFS file systems on the VHs.

```
# umount -a -t scatefs
```

6. Stop the scatefs-client service on the VHs.

```
# systemctl stop scatefs-client
```

7. Check the currently installed ScaTeFS package group ID (underlined part).

```
# yum group list installed -v 'scatefs-client-*'
(...)
Installed Groups:
   ScaTeFS Client for SX-Aurora TSUBASA and Mellanox OFED 4.9-0.1.7.0 (scatefs-client-tsubasa)
Done
```

   Uninstall the ScaTeFS client by specifying the group ID. The following is an example when the group ID is "scatefs-client-tsubasa".

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh scatefs-client-tsubasa
# yum group remove scatefs-client-tsubasa
```

⚠ Note

- If the file /opt/nec/ve/sbin/TSUBASA-groups-remark.sh does not exist, skip the execution of the script file TSUBASA-groups-remark.sh.
- If the error "No packages to remove from groups" occurs when executing yum group remove, execute yum groups mark remove as follows.

```
# yum group remove scatefs-client-tsubasa
No environment named scatefs-client-tsubasa exists
Maybe run: yum groups mark remove (see man yum)
No packages to remove from groups
# yum groups mark remove scatefs-client-tsubasa
Marked remove: scatefs-client-tsubasa
```

- If you do not have the PP support contract and cannot uninstall the ScaTeFS client specifying the group ID for it, uninstall it with the following command.

```
# yum remove '*scatefs-client-*'
```

8. Check that all ScaTeFS client packages are uninstalled.

```
# rpm -qa | grep scatefs-client
(Nothing displayed)
```

If any packages still remain, uninstall them as below:

```
# rpm -qa | grep scatefs-client
scatefs-client-xxxxx
# yum remove scatefs-client-xxxxx
```

5.3 Setup of the Yum Repository

You can use the NEC yum repository via the Internet or the one locally set up. Please refer to 3.1 for how to set up the local yum repository, if needed.

- In case of online

  Please execute the following command to install.

  - To set the default architecture to VE1

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

  - To set the default architecture to VE3

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

- In case of offline

  Please download the release package from the URL and store in the local yum repository for the free software.

- To set the default architecture to VE1

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

- To set the default architecture to VE3

```
# yum install https://sxauroratsubasa.sakura.ne.jp/repos/TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

Then, Please install the package as below.

- To set the default architecture to VE1

```
# cd /path/to/repos
# yum install ./TSUBASA-soft-release-ve1-3.0-1.noarch.rpm
```

- To set the default architecture to VE3

```
# cd /path/to/repos
# yum install ./TSUBASA-soft-release-ve3-3.0-1.noarch.rpm
```

⚠ Note

When the following messages are displayed during installation, it indicates that a configuration file was saved in a different file name (added .rpmnew extension) so that the new file doesn't overwrite an existing configuration. The new file includes configurations for new features. If user doesn't use new features, there is nothing to do about *.rpmnew file. If user will use new features, user needs to copy newly added configurations to the existing file and modify them for user's environment.

```
warning: /etc/yum.repos.d/TSUBASA-repo.repo created as /etc/yum.repos.d/TSUBASA-re
po.repo.rpmnew
warning: /etc/yum.repos.d/TSUBASA-restricted.repo created as /etc/yum.repos.d/TSUB
ASA-restricted.repo.rpmnew
```

Please refer to 3.1 for how to set up the local yum repository, if needed.

5.4 Stop of Update with the Yum Repository

If you are going to update the OS, disable update with the yum command by setting the enabled field to 0 in the following configuration files.

- /etc/yum.repos.d/TSUBASA-restricted.repo for the repository with restricted access
- /etc/yum.repos.d/TSUBASA-repo.repo for the repository with free access

5.5 Uninstallation of MLNX_OFED (Optional)

If you are going to update the OS, kernel, or MLNX_OFED and MLNX_OFED is installed, it is necessary to uninstall MLNX_OFED before the update.

Firstly, make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
```

Secondly, uninstall the software that depends on MLNX_OFED.

```
# yum remove 'libibverbs-ve*' 'libmlx5-ve*' 'libveib*' 'kmod-ve_peermem*' 'libibumad-ve*'
'sharp-ve*' ve-memory-mapping 'libvedma-ve*' 'rdma-core-ve*' ve_peermem 'libvesharp*'
```

Then uninstall MLNX_OFED.

```
# /usr/sbin/ofed_uninstall.sh
```

(Ref.) Please also refer to documents on the NVIDIA official home page.
https://docs.nvidia.com/networking/software/adapter-software/index.html

If RHEL/CentOS 8.4 is installed on the VH, please install libibverbs after uninstallation of MLNX_OFED.

```
# yum install libibverbs
```

5.6 Update of the OS (Optional)

Update the OS to the supported one.
The NEC support portal below lists the operating systems and their kernel versions verified for the SX-Aurora TSUBASA.

[SX-Aurora TSUBASA]Verified Linux kernel
https://www.support.nec.co.jp/en/View.aspx?id=4140100078

5.7 Update of the Kernel (Optional)

Update the kernel on the VHs to the supported version, and reboot the VHs.
The NEC support portal below lists the operating systems and their kernel versions verified for the SX-Aurora TSUBASA.

[SX-Aurora TSUBASA]Verified Linux kernel
https://www.support.nec.co.jp/en/View.aspx?id=4140100078

After the update, to avoid kernel update to a version that is not verified, please configure the yum command using the file /etc/yum.conf so that kernel packages are not updated. The following is an example of the description in the file /etc/yum.conf to avoid kernel update, where 'exclude=kernel*' is specified.

```
# vi /etc/yum.conf
[main]
exclude=kernel*
```

⚠ Note

Regarding the kernel version of the host machine, we recommend that you update the kernel to the latest version among the kernels that have been verified for operation in principle. If the latest version is not used, ve_drv module in ve_drv-kmod package may not be loaded and VE may not be ONLINE after installing Aurora software. In that case, you will need to downgrade vp-kmod and ve_drv-kmod packages. Please refer to "5.13 Confirmation of VE Driver Compatible with Linux kernel" for details.

5.8 Installation of MLNX_OFED (Optional)

If you use InfiniBand with the SX-Aurora TSUBASA and have uninstalled MLNX_OFED, please refer to "2.3 Installation of MLNX_OFED (Optional)" and install it onto the VHs again.

If you use NVIDIA SHARP and the version of SHARP is changed 2.6.1 or before to 2.7.0 or later, please refer to "4.8.4 NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)" and configure the firewall on each VHs.

5.9 Update of the Yum Repository

When you have updated the OS in Section 5.5, it is also necessary to update the yum repository. If the NEC yum repository on the Internet is available, update the version number in the "baseurl=" line to the one corresponding to the OS version in the configuration files for the yum repository (/etc/yum.repos.d/TSUBASA-restricted.repo for the paid software and /etc/yum.reops.d/TSUBASA-repo.repo for the free software).

Also, enable update with the yum command by re-setting the enabled field to 1 in the configuration files.

5.10 Uninstallation of Unnecessary Software

- Uninstallation of the ve-memory mapping packages

  If the VHs are not equipped with InfiniBand HCAs, uninstall the ve-memory-mapping package because update without the uninstallation will fail due to software dependency.

  Firstly, make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
```

  Then uninstall the ve-memory-mapping package.

```
# yum remove ve-memory-mapping
```

5.11 Update of the SX-Aurora TSUBASA Software

Update the SX-Aurora TSUBASA software with the value of the shell variable TSUBASA_GROUPS set to the group names to be updated according to whether the SX-Aurora TSUBASA has InfiniBand and which paid software you have purchased. The value of the shell variable TSUBASA_GROUPS must not include the group names for the software you do not want to update. Please refer to Section 3.2.4 for the available group names.

⚠ Note

- The value of the shell variable TSUBASA_GROUPS should not include "scatefs-client-tsubasa", because the ScaTeFS client should be installed in Section 5.12.
- See also appendix to install multiple instance of SDK(Compiler, Numeric Library Collection, and NLCPy) and MPI.
- VEOS v3.0.2 improves Process accounting feature and changes the accounting file format. The psacct-ve package converts file format of uncompressed existing accounting files into new format during its installation. The existing

files are saved in the same directory.

⚠ Note

If you use NEC MPI 2.x before update, change "mpi_el8" in the baseurl described in the repository configuration file to "mpi_mofed5_el8".

The repository configuration file and repository name that need to be changed are as follows.

Configuration file: /etc/yum.repos.d/TSUBASA-repo.repo
Repository name: nec-mpi-runtime, nec-mpi-community

When you use SDK-SE or Compat C++, in addition to the above, change the repository of the following repository configuration file.

Configuration file：/etc/yum.repos.d/TSUBASA-restricted.repo
Repository name：nec-mpi

Example:
/etc/yum.repos.d/TSUBASA-repo.repo

```
[nec-mpi-runtime]
…
baseurl=https://sxauroratsubasa.sakura.ne.jp/repos/runtime/mpi/mpi_mofed5_el8



[nec-mpi-community]

…
baseurl=https://sxauroratsubasa.sakura.ne.jp/repos/community/mpi/mpi_mofed5_
el8
```

- for the environment to compile and execute programs

  Make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

  ```
  # /opt/nec/ve/sbin/terminate-all-ve-services
  ```

  Then, update Aurora Software.

```
# TSUBASA_GROUPS="ve-devel ve-infiniband nec-sdk-devel nec-mpi-devel nqsv-execution"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

Note that the compilers (nfort, ncc, and nc++) invoked from the path /opt/nec/ve/bin/ become the updated versions with the commands above. If you do not want to change the versions of the compilers invoked from the path /opt/nec/ve/bin/, please refer to Appendix A.1.3.

Execute the following script to start the services for the SX-Aurora TSUBASA after the update.

```
# /opt/nec/ve/sbin/start-all-ve-services
```

- for the environment only to execute programs

Make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
```

Then, update Aurora Software

```
# TSUBASA_GROUPS="ve-runtime ve-infiniband nec-sdk-runtime nec-mpi-runtime nqsv-execution"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

Execute the following script to start the services for the SX-Aurora TSUBASA after the update.

```
# /opt/nec/ve/sbin/start-all-ve-services
```

- for the frontend machine

```
# TSUBASA_GROUPS="ve-frontend nec-sdk-frontend nec-mpi-frontend"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

Note that the compilers (nfort, ncc, and nc++) invoked from the path /opt/nec/ve/bin/ become the updated versions with the commands above. If you do not want to change the versions of the compilers invoked from the path /opt/nec/ve/bin/, please refer to Appendix A.1.3.

5.12 Installation of the ScaTeFS Client (Optional)

If you have uninstalled the ScaTeFS client in Section 5.2, install it again according to the following procedure, which installs the latest version of the ScaTeFS client placed in the yum repository.

1. Install the ScaTeFS client again.

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh scatefs-client-tsubasa
# yum group install scatefs-client-tsubasa
```

⚠ Note

- If you do not have the PP support contract and cannot install the ScaTeFS client specifying the group ID for it, prepare the zip file of the ScaTeFS client packages downloaded in the internet delivery product download service, and install them. For the procedures, refer to section "3.2.3 Without the paid software or PP Support".

2. Restore the configuration file for the ScaTeFS client if the backup file /etc/scatefs/client/ibdevice.conf.rpmsave has been created at the uninstallation. If this file does not exist, skip this procedure.

```
# mv /etc/scatefs/client/ibdevice.conf.rpmsave /etc/scatefs/client/ibdevice.conf
mv: overwrite '/etc/scatefs/client/ibdevice.conf'? y
```

5.13 Confirmation of VE Driver Compatible with Linux kernel

Downgrading VE driver (vp-kmod and ve_drv-kmod packages) is required when the updated VE driver is not compatible with Linux kernel in host machine.
Regarding the kernel version of the host machine, we recommend that you update the kernel to the latest version among the kernels that have been verified for operation in principle (2.2 Kernel Update). If the latest version is not used, ve_drv module in ve_drv-kmod package may not be loaded and VE may not be ONLINE after installing Aurora software. In that case, you will need to downgrade vp-kmod and ve_drv-kmod packages.

5.13.1 Downgrade Confirmation

Check to see if downgrades are needed. First, check to see if the ve_drv module is loaded.

```
# lsmod | grep ve_drv
ve_drv              262144  128
vp                  20480  1 ve_drv
```

If ve_drv is displayed, ve_drv is loaded. Please go to section 5.14.

```
# lsmod | grep ve_drv
(Nothing displayed)
```

If nothing is displayed, it has not been loaded. Check the Linux kernel version of VH with the following command.

```
# uname -r
4.18.0-553.8.1.el8_10.x86_64
```

Next, check the versions of the vp-kmod and ve_drv-kmod packages and make sure that the coreesponding packages are installed for your kernel.

[SX-Aurora TSUBASA]Verified Linux kernel

https://www.support.nec.co.jp/en/View.aspx?id=4140100078

```
# rpm -q vp-kmod
vp-kmod-3.0.1-4.18.0_372.32.1_1.el8.x86_64
# rpm -q ve_drv-kmod
ve_drv-kmod-3.0.1-4.18.0_372.32.1.el8_6.x86_64_1.el8.x86_64
```

If the version of the installed packages is newer than the version of the vp-kmod and ve_drv-kmod packages corresponding to the Linux kernel version you are using, you need to downgrade the vp-kmod and ve_drv-kmod packages.

5.13.2 Downgrade of the Latest vp-kmod and ve_drv-kmod Packages

Downgrade the vp-kmod and ve_drv-kmod packages by executing the following commands.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
# yum downgrade vp-kmod ve_drv-kmod
# /opt/nec/ve/sbin/start-all-ve-services
```

5.14 Confirm VE device path information file

If your model is defined on Table 3 (in 1.2.1 Hardware), please confirm the VE device path information file.

Confirm the VE device path information file using the following command:

```
# ls -l /etc/opt/nec/ve/veos/ve_nodes.json
-rw-r--r--. 1 root root 274  Feb 22  2023 /etc/opt/nec/ve/veos/ve_nodes.json
```

If the VE device path information file exists. Please go to section 5.15.

If VE device path information file does not exist, execute the following command:

```
# /opt/nec/ve/mmm/analysis/sbin/mmm-config-json.sh 2
```

After executing the above commands, reconfirm the VE device path information file and then execute the following restart ve-services commands:

```
# ls -l /etc/opt/nec/ve/veos/ve_nodes.json
-rw-r--r--. 1 root root 274  Feb 22  2023 /etc/opt/nec/ve/veos/ve_nodes.json
# /opt/nec/ve/sbin/terminate-all-ve-services
# /opt/nec/ve/sbin/start-all-ve-services
```

5.15 Status Check of the VEs

Confirm the status of the VEs is ONLINE with the vecmd command as the superuser. It can take a few minutes until it becomes ONLINE.

```
# /opt/nec/ve/bin/vecmd state get
Vector Engine MMM-Command v1.0.0
Command:
state -N 0 get
--------------------------------------------------------------------------------
VE0 [03:00.0] [ ONLINE ] Last Modif:2017/11/29 10:18:00
--------------------------------------------------------------------------------
Result: Success
```

⚠ Note

If "UNINITIALIZED" or "OFFLINE" is displayed, please wait for a short while.
If VE state does not become ONLINE after waiting, please restart services by executing the below commands.
 # /opt/nec/ve/sbin/terminate-all-ve-services
 # /opt/nec/ve/sbin/start-all-ve-services
And if "There is no executable ve card!" is displayed, please restart services in the same way.

If "UNAVAILABLE" is displayed, please restart the VHs by reboot command.

```
# reboot
```

If "UNAVAILABLE" is still displayed after restarting, there may be a hardware failure. Please refer to the following guide and perform the separation.

- ✓ **NEC Aurora Forum Documentation**
  "Vector Engine 2.0 Troubleshooting Guide"

5.16 Update of the VMC Firmware

Check whether update of the VMC firmware is required as follows:

```
# /opt/nec/ve/bin/vecmd fwup check
```

If the message "Updating VMCFW is required." is displayed, you need to update the VMC firmware. In this case, perform the following operations as the root user.

(1) Change the status of the VEs and update the firmware by vecmd commands

- When updating to VE1/VE2

```
# /opt/nec/ve/bin/vecmd state set off
# /opt/nec/ve/bin/vecmd state set mnt
# /opt/nec/ve/bin/vecmd fwup vmc aurora_MK10.bin
```

* The update of the firmware will take about a few minutes.

Reboot the VHs

```
# reboot
```

- When updating to VE3

```
# /opt/nec/ve/bin/vecmd state set off
# /opt/nec/ve/bin/vecmd fwup
```

* The update of the firmware will take about 10 minutes.

Shutdown the VH and then turn it on.

Please note that this is different from VE1.

```
#shutdown -h now
```

(2) Log in as the root user

(3) Check the Status of the VEs

Confirm the status of the VEs is ONLINE. It can take a few minutes until it becomes

ONLINE.

```
# /opt/nec/ve/bin/vecmd state get
Vector Engine MMM-Command v1.0.0
Command:
state -N 0 get
--------------------------------------------------------------------------------
VE0 [03:00.0] [ ONLINE ] Last Modif:2017/11/29 10:18:00
--------------------------------------------------------------------------------
Result: Success
```

⚠ Note

If "UNINITIALIZED" or "OFFLINE" is displayed, please wait for a short while.
And, if some VEs are not ONLINE or not appear, please shutdown VH and power on VH.

5.17 Command Compatible Packages Installation

Starting with the March 2023 software release of SX-Aurora TSUBASA, commands ported from Linux will be prefixed with "ve-". If you need compability with commands ported from Linux before March 2023, please install the command compability package.

- VH for compiling and executing program

```
# yum install ve-command-name-compat-devel
```

- VH for executing program

```
# yum install ve-command-name-compat-runtime
```

- Frontend machine

```
# yum install ve-command-name-compat-front-end
```

5.18 Start of the ScaTeFS Client (Optional)

If you use the ScaTeFS client, start it by rebooting the VHs. If you have rebooted them at the update of the VMC firmware, this operation is not needed.

```
# reboot
```

5.19 Update of Software Configuration

After software update, the required resource might be changed and new features

might become available. Please refer to the following sections and update the configuration.

- 4.8.5 Configuration of persistent device name for HCAs

  From the end of October 2024 release, this setting has become necessary. For systems using InfiniBand, please configure the setting if it has not been set yet. After configuring, refer to "4.8.6 Binding HCAs to VEs and Selecting Communication Method (Optional)" and modify the contents of necmpi.conf.

- 4.12 Configuration of HugePages

  HugePages is configured automatically since the release of September 2021. To use the functionality, please comment out options for Hugepages in the configuration file for sysctl (e.g. /etc/sysctl.conf).

```
# vi /etc/sysctl.conf
## vm.nr_hugepages = 384  # comment out
```

- 4.13 Configuration for Partial Process Swapping

5.20 Start of the System Operation

5.20.1 Status Check of the VEs

Confirm the status of the VEs on the VHs targeted for the update is ONLINE with the vecmd command as the superuser. It can take a few minutes until it becomes ONLINE.

```
# /opt/nec/ve/bin/vecmd state get
Vector Engine MMM-Command v1.0.0
Command:
state -N 0 get
--------------------------------------------------------------------------------
VE0 [03:00.0] [ ONLINE ] Last Modif:2017/11/29 10:18:00
--------------------------------------------------------------------------------
Result: Success
```

⚠ Note

If "UNINITIALIZED" or "OFFLINE" is displayed, please wait for a short while.

5.20.2 Start of Monitoring of the VHs

If you use the monitoring software (Zabbix or Ganglia+Nagios) for the VHs, please bring the VHs back to the monitoring mode from the maintenance mode.

5.20.3 Addition of the VHs to the Job Scheduler

If you use the job scheduler, add the VHs targeted for the update to the job scheduler according to the following procedure.

1. Start the job server and launcher on the VHs with the systemctl command

```
# systemctl start nqs-jsv
```

2. Bind the job server on the VHs to the queue to be operated with the bind execution_queue subcommand of the qmgr(1M) command on the host where NQSV/Client is installed. For example, the job server with job server ID 100 is bound to the bq queue by the following operation.

```
$ qmgr -Pm
Mgr: bind execution_queue job_server bq job_server_id = 100
```

   Use the bind interactive_queue subcommand for interactive queues.

Chapter6 Uninstallation

This section explains how to uninstall the SX-Aurora TSUBASA software from VHs.

6.1 Removal of VHs from System Operation

6.1.1 Disconnection of VHs from the Job Scheduler

If you use the job scheduler, disconnect the VHs targeted for uninstallation from the job scheduler.

1. Unbind the job server on the VHs from the queue which is being operated using the unbind execution_queue subcommand of the qmgr(1M) command on the host where NQSV/Client is installed. For example, the job server with job server ID 100 is unbound from the bq queue by the following operation.

```
$ qmgr -Pm
Mgr: unbind execution_queue job_server bq job_server_id = 100
```

   Use the unbind interactive_queue subcommand for interactive queues.

2. Make sure that the sstat(1) command with the -J option does not display the job server ID unbound in Step 1 in the JSVNO field, to confirm that no jobs exist on the job server.

```
$ qstat -J -Pm
```

   If any jobs exist on the job server, perform one of the following operations according to your operation policy.

   - Wait for the jobs to complete execution.
   - Re-run the jobs on different hosts with the qrerun command.
   - Delete the jobs with the qdel command.

3. Stop the job server and launcher on the VHs using the systemctl command.

```
# systemctl stop nqs-jsv.target
```

6.1.2 Removal of the VHs from Target of Monitoring

If you use the monitoring software (Zabbix or Ganglia+Nagios) for the VHs, please place the VHs into the maintenance mode.

6.2 Uninstallation of the ScaTeFS Client (Optional)

If you use the ScaTeFS client on the VHs, stop the ScaTeFS client service and uninstall the ScaTeFS client according to the following procedure.

1. If a ScaTeFS file system on the VHs targeted for uninstallation is exported via NFS, unmount the ScaTeFS file system on all the NFS clients, then go to Step 2. Otherwise, go to Step 3.

2. Stop the nfs-server service on the VHs.

```
# systemctl stop nfs-server
```

3. If a ScaTeFS file system on the VHs is exported as a share via Samba, stop accessing the Samba share on all the CIFS clients such as Windows, then go to Step 4.
   Otherwise, go to Step 5.

4. Stop the smb service and nmb service on the VHs.

```
# systemctl stop smb
# systemctl stop nmb
```

5. Unmount the ScaTeFS file systems on the VHs.

```
# umount -a -t scatefs
```

6. Stop the scatefs-client service on the VHs.

```
# systemctl stop scatefs-client
```

7. Check the currently installed ScaTeFS package group ID (underlined part).

```
# yum group list installed -v 'scatefs-client-*'
(...)
Installed Groups:
   ScaTeFS Client for SX-Aurora TSUBASA and Mellanox OFED 4.9-0.1.7.0 (scatefs-client-tsubasa)
Done
```

Uninstall the ScaTeFS client by specifying the group ID. The following is an example when the group ID is "scatefs-client-tsubasa".

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh scatefs-client-tsubasa
# yum group remove scatefs-client-tsubasa
```

⚠ Note

- If the file /opt/nec/ve/sbin/TSUBASA-groups-remark.sh does not exist, skip the execution of the script file TSUBASA-groups-remark.sh.
- If the error "No packages to remove from groups" occurs when executing yum group remove, execute yum groups mark remove as follows.

```
# yum group remove scatefs-client-tsubasa
No environment named scatefs-client-tsubasa exists
Maybe run: yum groups mark remove (see man yum)
No packages to remove from groups
# yum groups mark remove scatefs-client-tsubasa
Marked remove: scatefs-client-tsubasa
```

- If you do not have the PP support contract and cannot uninstall the ScaTeFS client specifying the group ID for it, uninstall it with the following command.

```
# yum remove '*scatefs-client-*'
```

8. Check that all ScaTeFS client packages are uninstalled.

```
# rpm -qa | grep scatefs-client
(Nothing displayed)
```

If any packages still remain, uninstall them as below:

```
# rpm -qa | grep scatefs-client
scatefs-client-xxxxx
# yum remove scatefs-client-xxxxx
```

6.3 Uninstallation

Make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
```

Uninstall the SX-Aurora TSUBASA software with the yum command.

```
# TSUBASA_GROUPS="ve-devel ve-infiniband nec-sdk-devel nec-mpi-devel nqsv-execution ve-runtime nec-sdk-runtime nec-mpi-runtime ve-frontend"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
```

```
# yum group remove $TSUBASA_GROUPS
# yum remove TSUBASA-soft-release-ve[13]
```

Confirm that the SX-Aurora TSUBASA software is not installed with the following command. If some SX-Aurora TSUBASA software remains, uninstall it with the yum remove command. Please note that it is not necessary to uninstall NEC software other than the SX-Aurora TSUBASA software.

```
# rpm -qa --qf '%{VENDOR} %{NAME}-%{VERSION}-%{RELEASE}.%{ARCH}\n'| grep NEC
```

Appendix A Installation of the Software Supporting Multiple Instances

A.1 SDK（Compilers）

In this section, alternative commands are delimited by a vertical bar and enclosed by square brackets. For instance, /path/[nfort|ncc|nc++] means /path/nfort, /path/ncc, or /path/nc++.

A.1.1 Installation of a Specific Version of the Compilers

The following examples show how to install a specific version of the compilers. "X.X.X" in the examples corresponds to the compiler version.

- Example 1: for the environment to compile and execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-compiler-devel-X.X.X
# yum group install nec-compiler-devel-X.X.X
```

- Example 2: for the environment only to execute programs

  It is not necessary to install a specific version of the compilers because the latest versions also include older versions. Update the compilers to the latest versions referring to Section 5.11.

- Example 3: On the frontend machines

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-compiler-frontend-X.X.X
# yum group install nec-compiler-frontend-X.X.X
```

The compiler version invoked with the command /opt/nec/ve/bin/[nfort|ncc|nc++] is not changed by the procedure in this subsection. Please specify the compiler version explicitly after the compiler name in command lines to invoke the compilers installed in this subsection. For example, the nfort compiler installed here is invoked with the command /opt/nec/ve/bin/nfort-X.X.X, where X.X.X is the compiler version. Please refer to the next section for how to change the compiler version invoked with the command /opt/nec/ve/bin/[nfort|ncc|nc++].

A.1.2 Change of the Compiler Versions Invoked with the Command /opt/nec/ve/bin/[nfort|ncc|nc++]

It is possible to change the compiler versions (nfort, ncc, and nc++) invoked with the command /opt/nec/ve/bin/[nfort|ncc|nc++] using the yum command as follows.

- Example: for the environment to compile and execute programs

```
# yum remove nec-nfort-inst
# yum remove nec-nc++-inst
# yum install nec-nfort-inst-X.X.X-X.noarch
# yum install nec-nc++-inst-X.X.X-X.noarch
```

A.1.3 Update of the Compilers without Changing the Versions Invoked with the Command /opt/nec/ve/bin/[nfort|ncc|nc++]

It is possible to update the compilers (nfort, ncc, and nc++) without changing the compiler versions invoked with the command /opt/nec/ve/bin/[nfort|ncc|nc++] by specifying the shell variable TSUBASA_GROUPS as follows.

- Example1: for the environment to compile and execute programs

  Make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
```

  Then, update Aurora Software.

```
# TSUBASA_GROUPS="ve-devel ve-infiniband nec-sdk-devel-nodefault nec-mpi-devel nqsv-execution"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

  Execute the following script to start the services for the SX-Aurora TSUBASA after the update.

```
# /opt/nec/ve/sbin/start-all-ve-services
```

- Example2: for the environment only to execute programs

  Make sure that any programs are not being executed on the VEs, and then execute the following script to stop the services for the SX-Aurora TSUBASA.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
```

Then, update Aurora Software

```
# TSUBASA_GROUPS="ve-runtime ve-infiniband nec-sdk-runtime nec-mpi-runtime nqsv-execution"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

Execute the following script to start the services for the SX-Aurora TSUBASA after the update.

```
# /opt/nec/ve/sbin/start-all-ve-services
```

- Example3: for the frontend machine

```
# TSUBASA_GROUPS="ve-frontend nec-sdk-frontend-nodefault nec-mpi-frontend"
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh $TSUBASA_GROUPS
# yum group install $TSUBASA_GROUPS
```

The updated version of the nfort compiler is invoked with the command /opt/nec/ve/bin/nfort-X.X.X, where X.X.X is the compiler version.

- Example: Invocation of nfort version 2.4.1

```
$ /opt/nec/ve/bin/nfort-2.4.1 t.f90
```

A.2 MPI

The following examples show how to install a specific version of the MPI. "X-X-X" in the examples corresponds to the MPI Version X.X.X. Please check the following note before installation.

- Example 1: for the environment to compile and execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-mpi-devel-X-X-X
# yum group install nec-mpi-devel-X-X-X
```

- Example 2: for the environment only to execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-mpi-runtime-X-X-X
# yum group install nec-mpi-runtime-X-X-X
```

- Example 3: On the frontend machines

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-mpi-frontend-X-X-X
# yum group install nec-mpi-frontend-X-X-X
```

⚠ Note

Commands such as mpirun placed in /opt/nec/ve/bin are not updated version-specified installation. To update the command placed in /opt/nec/ve/bin, update the nec-mpi-runtime package to the specified version with the following command.

```
# yum install nec-mpi-runtime-X.X.X
```

When updating to the latest version including the command placed in /opt/nec/ve/bin, perform the installation without specifying the version.

If you perform the installation specifying a version, the nec-mpi-runtime-common package is required. Please complete the installation of the nec-mpi-runtime-common package with the following command before staring work.

```
# yum install nec-mpi-runtime-common
```

A.3 Numeric Library Collection

The following examples show how to install a specific version of Numeric Library Collection. "X.X.X" in the examples corresponds to the version of Numeric Library Collection.

- Example 1: for the environment to compile and execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-nlc-devel-X.X.X
# yum group install nec-nlc-devel-X.X.X
```

- Example 2: for the environment only to execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-nlc-runtime-X.X.X
# yum group install nec-nlc-runtime-X.X.X
```

- Example 3: On the frontend machines

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-nlc-frontend-X.X.X
# yum group install nec-nlc-frontend-X.X.X
```

A.4 NLCPy

The following examples show how to install a specific version of NLCPy. "X.X.X" in the examples corresponds to the version of NLCPy.

- Example 1: for the environment to compile and execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-nlcpy-devel-X.X.X
# yum group install nec-nlcpy-devel-X.X.X
```

- Example 2: for the environment only to execute programs

```
# /opt/nec/ve/sbin/TSUBASA-groups-remark.sh nec-nlcpy-runtime-X.X.X
# yum group install nec-nlcpy-runtime-X.X.X
```

Appendix B Installed packages

B.1 SX-Aurora Software Packages

Table 13 The List of the SX-Aurora TSUBASA Software

Software	Package File	How to get (*)
VEOS	autoconf-ve automake-ve coreutils-ve gdb-ve glibc-ve1 glibc-ve1-devel glibc-ve3 glibc-ve3-devel glibc-ve-wrapper kheaders-ve1 kheaders-ve3 libgcc-ve1-static libgcc-ve3-static libsysve-ve1 libsysve-ve1-devel libsysve-ve3 libsysve-ve3-devel libthread_db-ve1 libthread_db-ve3 libtool-ve libved procps-ng-ve psacct-ve psmisc-ve strace-ve sysstat-ve time-ve util-linux-ve ve_drv-kmod vedebuginfo velayout venumainfo veoffload-aveo veoffload-aveo-devel veoffload-aveorun-ve1 veoffload-aveorun-ve1-devel veoffload-aveorun-ve3 veoffload-aveorun-ve3-devel veoffload-veda veoffload-veda-devel veoffload-veda-ve1 veoffload-veda-ve3 veos veos-devel veos-headers veosinfo veosinfo3 veos-libveptrace	A

Software	Package File	How to get (*)
	vesysinit vesysinit-udev vp-kmod veswap libvhcall-fortran-ve1 libvhcall-fortran-ve1-devel libvhcall-fortran-ve3 libvhcall-fortran-ve3-devel libsmartcols-ve cmake-ve veosctl ve-command-name-compat-devel ve-command-name-compat-front-end ve-command-name-compat-runtime	
InfiniBand for SX-Aurora TSUBASA	libveib-mofedX.X kmod-ve_peermem-mofedX.X rdma-core-ve1-mofedX.X libibverbs-ve1-mofedX.X libibumad-ve1-mofedX.X sharp-ve1-X.X.X libvesharp-X.X.X rdma-core-ve3-mofedX.X libibverbs-ve3-mofedX.X libibumad-ve3-mofedX.X sharp-ve3-X.X.X	A
MMM	ftmon ftmon-ve3 mmm mmm-config mmm-analysis mmm-analysis-ve3 mmm-msl mmm-msl-ve3 rtmon rtmon-ve3 ve-firmware	A
VMC Firmware	vmcfw vmcfw-ve3	A
License Access Library	aurlic-lib	A
ScaTeFS Client	kmod-scatefs-client-modules-mofedX.X scatefs-client-libscatefsib_ve scatefs-client-mount-utils scatefs-client-rcli-utils scatefs-client-sys-utils scatefs-client-usr-utils	B(ScaTeFS/Client)

Software	Package File	How to get (*)
NEC MPI	nec-mpi-devel-X-X-X nec-mpi-devel-ve3-X-X-X nec-mpi-libs-X-X-X nec-mpi-libs-ve3-X-X-X nec-mpi-utils-X-X-X nec-mpi-utils-ve3-X-X-X nec-mpi-runtime-X-X-X nec-mpi-runtime-ve3-X-X-X nec-mpi-runtime nec-mpi-runtime-common nec-mpi-runtime-pbs	Both A and B
Tuning Tool	nec-veperf-bin nec-veperf-ve1-devel nec-veperf-ve3-devel nec-veperf-ve1-libs nec-veperf-ve3-libs nec-ftraceviewer	Both A and B
NEC Parallel Debugger	nec-paralleldebugger	B
NQSV/JobServer	NQSV-JobServer	B(NQSV/JobServer)
NQSV/Client	NQSV-Client	B(NQSV/Resource Manager)
Numeric Library Collection	nec-asl-ve-X.X.X nec-aslfftw-ve-X.X.X nec-blas-ve-X.X.X nec-heterosolver-ve-X.X.X nec-lapack-ve-X.X.X nec-nlc-base-X.X.X nec-nlc-doc-X.X.X nec-sblas-ve-X.X.X nec-sca-ve-X.X.X nec-scalapack-ve-X.X.X nec-asl-ve-devel-X.X.X nec-aslfftw-ve-devel-X.X.X nec-blas-ve-devel-X.X.X nec-heterosolver-ve-devel-X.X.X nec-lapack-ve-devel-X.X.X nec-sblas-ve-devel-X.X.X nec-sca-ve-devel-X.X.X nec-scalapack-ve-devel-X.X.X	Both A and B
NLCPy	nec-python36-nlcpy-ve-X.X.X nec-python36-numpy-X.X.X-Y.Y.Y	Both A and B
mpi4py-ve	nec-python36-mpi4py-ve-X.X.X	Both A and B
binutils	binutils-ve	Both A and B

Software	Package File	How to get (*)
C/C++ compiler	nec-nc++-X.X.X nec-nc++-inst nec-nc++-doc-X.X.X nec-nc++-shared-X.X.X nec-nc++-shared-devel-X.X.X nec-nc++-shared-inst-X nec-nc++-compat (*) nec-nc++-compat-shared nec-nc++-compat-shared-devel nec-nc++-compat-shared-inst	Both A and B (*) is only in B(SDK-SE)
Fortran compiler	nec-nfort-X.X.X nec-nfort-inst nec-nfort-doc-X.X.X nec-nfort-shared-X.X.X nec-nfort-shared-devel-X.X.X nec-nfort-shared-inst-X nec-nfort-runtime	Both A and B
NEC Compat C++ Standard Library for Vector Engine	nec-nc++-compat nec-nc++compat-shared nec-nc++-compat-shared-devel nec-nc++-compat-shared-inst	B

(*)A: Free software. You can install the software packages from the NEC yum repository with free access with the yum command.

B: Paid software. If you have the PP support contract, you can install the software packages from the NEC yum repository with restricted access with the yum command. Otherwise, you can install the packages downloaded in the internet delivery product download service.

B.2 Packages for SX-Aurora Software Package Group

The package files installed and updated in the Package Group described in Section 3.2.4 are as follows.

Table 14 Package Group/InfiniBand for SX-Aurora TSUBASA

Package	Group Name
	ve-infiniband
libibverbs-ve1-mofedX.X	✓
rdma-core-ve1-mofedX.X	✓
libveib-mofedX.X	✓
ve-memory-mapping	✓
kmod-ve_peermem-mofedX.X	✓
libibumad-ve1-mofedX.X	✓
sharp-ve1-X.X.X	✓
libvesharp-X.X.X	✓
libibverbs-ve3-mofedX.X	✓
rdma-core-ve3-mofedX.X	✓
libibumad-ve3-mofedX.X	✓
sharp-ve3-X.X.X	✓

Table 15 Package Group/VE Application

Package	Group Name		
	ve-devel	ve-runtime	ve-frontend
aurlic-lib	✓	✓	✓
aurlic-server	✓	✓	✓
autoconf-ve	✓	-	✓
automake-ve	✓	-	✓
coreutils-ve	✓	✓	-
ftmon	✓	✓	-
ftmon-ve3	✓	✓	-
gdb-ve	✓	✓	✓
glibc-ve1	✓	✓	✓
glibc-ve1-devel	✓	-	✓
glibc-ve3	✓	✓	✓
glibc-ve3-devel	✓	-	✓
glibc-ve-wrapper	✓	✓	✓

Package	Group Name		
	ve-devel	ve-runtime	ve-frontend
kheaders-ve1	✓	-	✓
kheaders-ve3	✓	-	✓
libgcc-ve1-static	✓	-	✓
libgcc-ve3-static	✓	-	✓
libsysve-ve1	✓	✓	✓
libsysve-ve1-devel	✓	-	✓
libsysve-ve3	✓	✓	✓
libsysve-ve3-devel	✓	-	✓
libthread_db-ve1	✓	✓	✓
libthread_db-ve3	✓	✓	✓
libtool-ve	✓	-	✓
libved	✓	✓	✓
mmm	✓	✓	-
mmm-analysis	✓	✓	-
mmm-analysis-ve3	✓	✓	-
mmm-config	✓	✓	-
mmm-msl	✓	✓	-
mmm-msl-ve3	✓	✓	-
procps-ng-ve	✓	✓	-
psacct-ve	✓	✓	-
psmisc-ve	✓	✓	-
rtmon	✓	✓	-
rtmon-ve3	✓	✓	-
strace-ve	✓	✓	-
sysstat-ve	✓	✓	-
time-ve	✓	✓	-
util-linux-ve	✓	✓	-
ve_drv-kmod	✓	✓	-
vedebuginfo	✓	-	✓
ve-firmware	✓	✓	-
velayout	✓	✓	✓
venumainfo	✓	✓	-

Package	Group Name		
	ve-devel	ve-runtime	ve-frontend
veoffload-aveo	✓	✓	-
veoffload-aveo-devel	✓	-	✓
veoffload-aveorun-ve1	✓	✓	✓
veoffload-aveorun-ve1-devel	✓	-	✓
veoffload-aveorun-ve3	✓	✓	✓
veoffload-aveorun-ve3-devel	✓	-	✓
veoffload-veda	✓	✓	✓
veoffload-veda-devel	✓	-	✓
veoffload-veda-ve1	✓	✓	✓
veoffload-veda-ve3	✓	✓	✓
veos	✓	✓	-
veos-devel	✓	-	✓
veos-headers	✓	-	✓
veosinfo	✓	✓	-
veosinfo3	✓	✓	-
veos-libveptrace	✓	✓	✓
vesysinit	✓	✓	-
vesysinit-udev	✓	✓	-
vmcfw	✓	✓	-
vmcfw-ve3	✓	✓	-
vp-kmod	✓	✓	-
veswap	✓	✓	-
libvhcall-fortran-ve1	✓	✓	✓
libvhcall-fortran-ve1-devel	✓	-	✓
libvhcall-fortran-ve3	✓	✓	✓

Package	Group Name		
	ve-devel	ve-runtime	ve-frontend
libvhcall-fortran-ve3-devel	✓	-	✓
libsmartcols-ve	✓	✓	-
cmake-ve	✓	-	✓
veosctl	✓	✓	-

*) The group ve-devel is used when both compilation and execution of the programs using VEs are performed on the target machine. ve-runtime is used when only execution of the programs using VEs is performed on the target machine. And ve-frontend is used for the frontend machines.

Table 16 Package Group/NEC SDK

Package	Group Name				
	nec-sdk-runtime	nec-sdk-devel (SDK-CE)	nec-sdk-frontend (SDK-CE)	nec-sdk-devel (SDK-SE/ Compat C++	nec-sdk-frontend (SDK-SE/ Compat C++)
binutils-ve	✓	✓	✓	✓	✓
nec-aslfftw-ve-X.X.X	✓	✓	✓	✓	✓
nec-aslfftw-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-asl-ve-X.X.X	✓	✓	✓	✓	✓
nec-asl-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-blas-ve-X.X.X	✓	✓	✓	✓	✓
nec-blas-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-ftraceviewer	✓(*2)	✓	✓	✓	✓
nec-heterosolver-ve-X.X.X	✓	✓	✓	✓	✓
nec-heterosolver-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-lapack-ve-X.X.X	✓	✓	✓	✓	✓
nec-lapack-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-python36-nlcpy-ve-X.X.X	✓	✓	-	✓	✓
nec-python36-numpy-X.X.X-Y.Y.Y	✓	✓	-	✓	✓

Package	Group Name				
	nec-sdk-runtime	nec-sdk-devel (SDK-CE)	nec-sdk-frontend (SDK-CE)	nec-sdk-devel (SDK-SE/ Compat C++	nec-sdk-frontend (SDK-SE/ Compat C++)
nec-python36-mpi4py-ve-X.X.X	✓	✓	-	✓	-
nec-nc++-X.X.X	-	✓	✓	✓	✓
nec-nc++-doc-X.X.X	-	✓	✓	✓	✓
nec-nc++-inst	-	✓	✓	✓	✓
nec-nc++-shared-X.X.X	✓	✓	✓	✓	✓
nec-nc++-shared-devel-X.X.X	-	✓	✓	✓	✓
nec-nc++-shared-inst-X	✓	✓	✓	✓	✓
nec-nc++-compat	-	-	-	✓	✓
nec-nc++-compat-shared-X.X.X	✓	✓	✓	✓	✓
nec-nc++-compat-devel-X.X.X	-	✓	✓	✓	✓
nec-nc++-compat-inst-X	✓	✓	✓	✓	✓
nec-nfort-X.X.X	-	✓	✓	✓	✓
nec-nfort-doc-X.X.X	-	✓	✓	✓	✓
nec-nfort-inst	-	✓	✓	✓	✓
nec-nfort-runtime	✓	✓	✓	✓	✓
nec-nfort-shared-X.X.X	✓	✓	✓	✓	✓
nec-nfort-shared-devel-X.X.X	-	✓	✓	✓	✓
nec-nfort-shared-inst-X	✓	✓	✓	✓	✓
nec-nlc-base-X.X.X	✓	✓	✓	✓	✓
nec-nlc-doc-X.X.X	-	✓	✓	✓	✓
nec-paralleldebugger	✓(*2)	-	-	✓	✓
nec-sblas-ve-X.X.X	✓	✓	✓	✓	✓
nec-sblas-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-scalapack-ve-X.X.X	✓	✓	✓	✓	✓
nec-scalapack-ve-devel-X.X.X	-	✓	✓	✓	✓

Package	Group Name				
	nec-sdk-runtime	nec-sdk-devel (SDK-CE)	nec-sdk-frontend (SDK-CE)	nec-sdk-devel (SDK-SE/ Compat C++	nec-sdk-frontend (SDK-SE/ Compat C++)
nec-sca-ve-X.X.X	✓	✓	✓	✓	✓
nec-sca-ve-devel-X.X.X	-	✓	✓	✓	✓
nec-veperf-bin	✓(*2)	✓	✓	✓	✓
nec-veperf-ve1-devel	-	✓	✓	✓	✓
nec-veperf-ve3-devel	-	✓	✓	✓	✓
nec-veperf-ve1-libs	✓	✓	✓	✓	✓
nec-veperf-ve3-libs	✓	✓	✓	✓	✓

(*1) The group nec-sdk-devel is used when both compilation and execution of the programs using VEs are performed on the target machine, nec-sdk-runtime is used when only execution of the programs using VEs is performed on the target machine, and nec-sdk-frontend is used for the frontend machines.

(*2) The packages are contained only in nec-sdk-runtime of the yum repository "SDK-SE/Compat C++".

Table 17 Package Group/NEC MPI

Package	Group Name		
	nec-mpi-runtime	nec-mpi-devel	nec-mpi-frontend
nec-mpi-devel-X-X-X	-	✓	✓
nec-mpi-devel-ve3-X-X-X	-	✓	✓
nec-mpi-libs-X-X-X	✓	✓	✓
nec-mpi-libs-ve3-X-X-X	✓	✓	✓
nec-mpi-utils-X-X-X	✓	✓	✓
nec-mpi-utils-ve3-X-X-X	✓	✓	✓
nec-mpi-runtime-X-X-X	✓	✓	✓
nec-mpi-runtime-ve3-X-X-X	✓	✓	✓
nec-mpi-runtime	✓	✓	✓
nec-mpi-runtime-common	✓	✓	✓
nec-mpi-runtime-pbs	✓	✓	✓

Appendix C Network Configuration

It is necessary to set up the network as follows to execute parallel programs using VEs on multiple VHs.

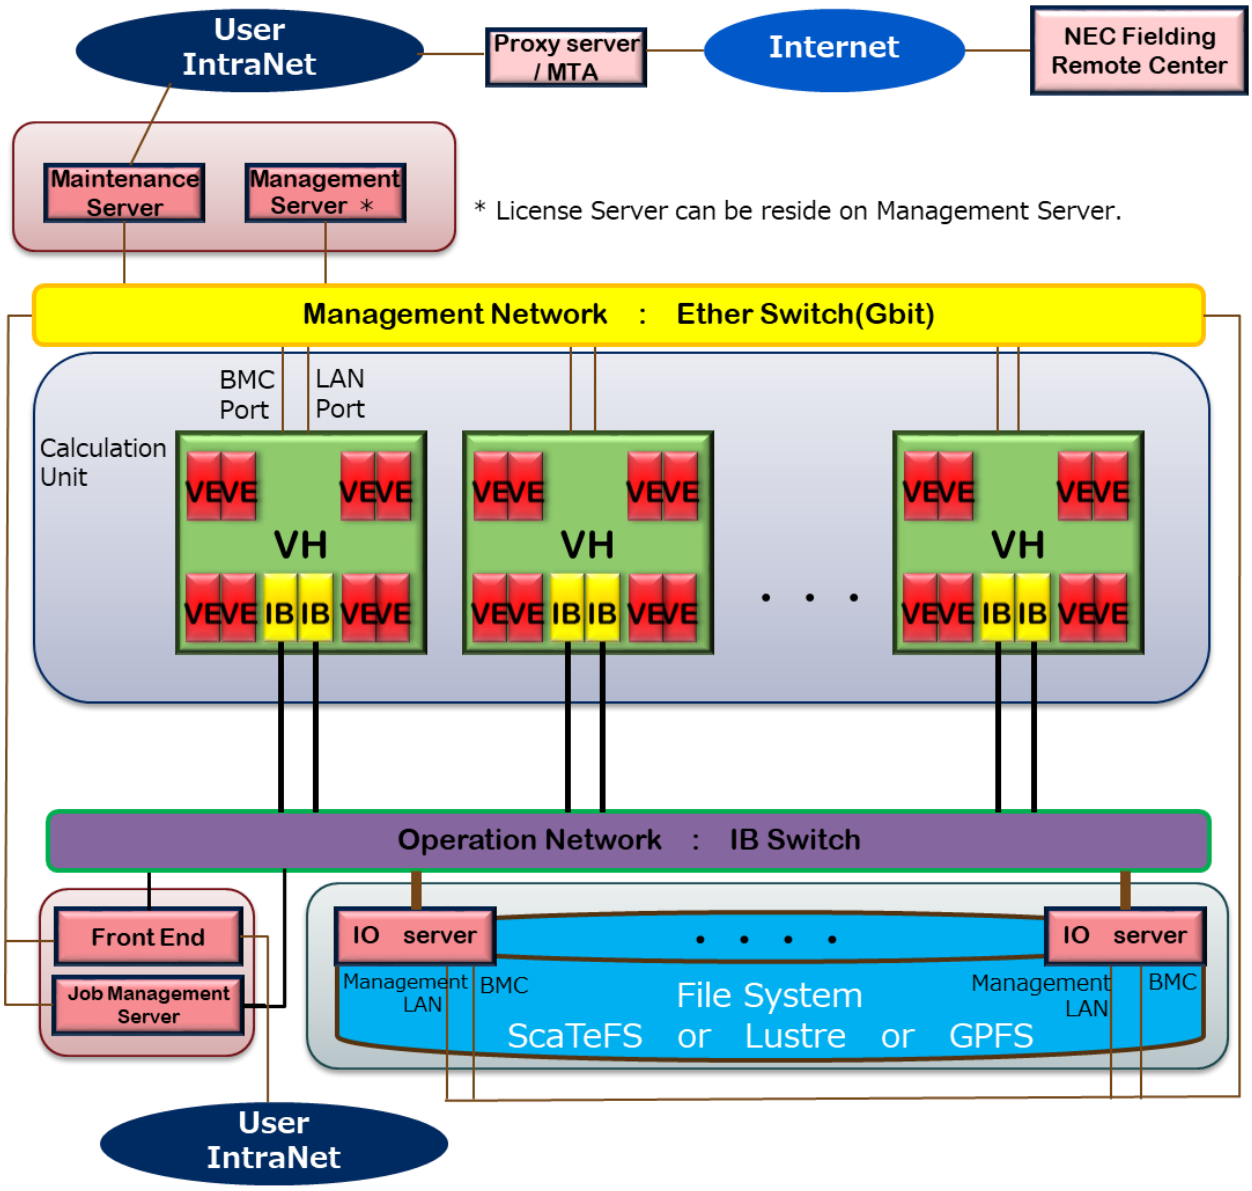

Figure 10 Network Configuration

C.1 Operation Network

The operation network has the following three functions.

a) MPI communication network (In the case of using InfiniBand)

b) I/O network (Network for file systems)

c) Operations management network (For job control and job transfer. The IP over IB interface is used in the case of using InfiniBand)

The operation network is a high bandwidth network where hosts are connected to InfiniBand switches with InfiniBand cables. Hosts in the network need to be equipped with InfiniBand Host Channel Adapters (HCA). It is not possible to separate the hosts physically, but possible to assign the maximum bandwidth to each function by the QoS setting. This network is mandatory to construct a VI cluster.
When using only single VI, you can construct the operation network with Ethernet.

Servers and devices expected to be connected to the network are shown below. It is possible to connect any servers and devices to the network according to requirements without limited to them.

a) Computing units (VIs, core part in the SX-Aurora TSUBASA system)
b) Frontend machines
c) Job management server
d) I/O servers (Connected to file systems)

C.2 Management Network

The management network has the following two functions.

a) Operations management network (System operation and monitoring of computational resources)
b) Maintenance management network (Monitoring of hardware)

The management network is connected with Ethernet. Gigabit Ethernet (GbE) is recommended. Hosts in the network are connected to Ethernet switches with LAN cables. VHs have LAN ports that can be used to connect to the network. It is strongly recommended that the BMC ports on VHs also be connected to the management network.
This document assumes that one of the LAN ports on each VH is connected to the management network.

Servers and devices expected to be connected to the network are shown below. It is

possible to connect any servers and devices to the network according to requirements without limited to them.

a) Computing units (VIs, core part in the SX-Aurora TSUBASA system)
b) Frontend machines
c) Job management server
d) I/O servers (and the management interface for storage)
e) Operations management server (including the license server)
f) Maintenance server
g) Yum repository server

Appendix D Migration to the Glibc Environment

The SX-Aurora TSUBASA software has supported the GNU C Library (glibc) as the official C library on VEs as of the release on January, 2019. The support of the former C library musl-libc was ceased at the end of March, 2019.

If you use the musl-libc environment now, please migrate to the glibc environment according to the following procedure.

1. Update the SX-Aurora TSUBASA software to the latest versions that support glibc.
2. Recompile your programs in the glibc environment.

Since musl-libc is not maintained anymore, please uninstall the musl-libc related packages by executing the following command.

```
# yum -y remove musl-libc-ve musl-libc-ve-devel libsysve-musl libsysve-musl-devel veos-musl-headers
```

There is no interoperability between binaries compiled with glibc and those with musl-libc. If you need to identify the C library linked to a binary file, execute the "ve-libc-check" script as follows. This script supports any kind of binaries for VEs such as "a.out", ".o", ".a", or ".so".

```
$ /opt/nec/ve/bin/ve-libc-check ./a.out
  This is compiled with musl-libc: /home/userxxx/a.out
```

The above output shows that the "a.out" file is compiled with musl-libc. If no message is printed, the binary is compiled with glibc and does not require musl-libc.

⚠ Note

- The support for musl-libc was ceased at the end of March 2019.
- The "ve-libc-check" script does not support object files created from ".s" files. Please be careful not to mix binaries compiled with musl-libc and those with glibc especially when you have ".s" source files.
- The "ve-libc-check" script does not support checking libraries dynamically linked to a program. For example, if a program compiled with and linked to glibc loads or links libraries compiled with and linked to musl-libc dynamically, the "ve-libc-check" script cannot check it. Please re-make all of your libraries with glibc.

Appendix E The manual setting of HugePages

(1) The Calculation of HugePages

The required number of HugePages can be calculated from Table 18. For example, suppose one VE with 8 cores. In this case, VEOS requires 32, NEC MPI requires 128. So, the number of HugePages should be set to the sum of these values, 160.

The size of HugePages is 2MB per page.

Table 18 The Required Number of HugePages

Package	The Required Number of HugePages
VEOS (Swap out to VH memory)	[The sum of the size of memory to be swapped out] / [2M] (*1)
VEOS (Swap out to files)	The number of VE * 32
VEOS (Others)	The number of VE cores * The number of VE * 4 (*2)
NEC MPI	The number of VE cores * The number of VE * 16 (*2)
Node with InfiniBand	+ The number of VH logical cores * 8 (*2) (*3)
Node without InfiniBand	+ The number of VH logical cores * 32 (*2) (*3)
ScaTeFS Client (2VE server)	The number of VE cores * The number of VE * 32 (*2)
ScaTeFS Client (4VE server)	The number of VE cores * The number of VE * 64 (*2)

*1 For example, suppose your machine has eight VEs and maximally 40GB VE memory is swapped out to VH memory per VE. In this case, 163,840(=8×40×1024/2) HugePages are required.

*2 When the Partial Process Swapping is valid, the number sould be double.

*3 Add the additional HugePages to the value of "NEC MPI".

If NQSV socket scheduling is required, refer to the NEC Network Queuing System V (NQSV) Users Guide [Administration], Chapter.18 Socket Scheduling, 18.1.3 Memory allocation policy.

There are cases in which Automatic HugePages configuration command allocates some HugePages based on VE architecture. In this case, there is possibility that the number of HugePages calculated using the Table 18 is larger than the one calculated by the command.

If user don't use the command, use the number of Hugepages calculated using the Table 18.

(2) Contiguration of HugePages

In this section, confiture the calculated number of Hugepages.

1. Check of the Current Number of HugePages

   Check the current number of HugePages with the following command. If the value is not zero, please make sure that the setting for the SX-Aurora software described in this section has not been made in the file /etc/sysctl.conf. When the number of HugePages has been set for some software other than the SX-Aurora TSUBASA software, please add it to the required number of HugePages calculated at (1).

```
# sysctl vm.nr_hugepages
vm.nr_hugepages = 0
```

2. Configuration of HugePages

   If the value of the parameter vm.nr_hugepages is already set in the file /etc/sysctl.conf, update the value with the one calculated in the previous step. Otherwise, add the line vm.nr_hugepages to the file to specify the value.

```
# vi /etc/sysctl.conf
vm.nr_hugepages = 384
```

   Apply the setting with the sysctl command.

```
# sysctl --system
```

3. Check of the Updated Number of HugePages

   Make sure using the sysctl command that the number of HugePages is updated with the calculated value ("vm.nr_hugepages = 384" in this example).

```
# sysctl vm.nr_hugepages
vm.nr_hugepages = 384
```

4. Restarting services

   Restart services in order to make services use huge pages.

```
# /opt/nec/ve/sbin/terminate-all-ve-services
# /opt/nec/ve/sbin/start-all-ve-services
```

5. Invalidate the automatic HugePages configuration services

Validate SKIP_SETTING option in /etc/opt/nec/ve/veos/ve-hugepages.conf in order to the automatic HugePages configuration command doesn't set a value.

```
# vi /etc/opt/nec/ve/veos/ve-hugepages.conf
SKIP_SETTING=YES
```

6. Status Check of the VEs

Confirm the status of the VEs is ONLINE with the vecmd command as the superuser. It can take a few minutes until it becomes ONLINE.

```
# /opt/nec/ve/bin/vecmd state get
Vector Engine MMM-Command v1.0.0
Command:
state -N 0 get
--------------------------------------------------------------------------------
VE0 [03:00.0][ ONLINE ] Last Modif:2017/11/29 10:18:00
--------------------------------------------------------------------------------
Result: Success
```

⚠ Note

If "UNINITIALIZED" or "OFFLINE" is displayed, please wait for a short while.

Appendix F History

History table

Feb. 2018	1st edition
Mar. 2023	48th edition
Jun. 2023	49th edition
Sep. 2023	50th edition
Oct. 2023	51st edition
Feb. 2024	52nd edition
Oct. 2024	53rd edition
Feb. 2025	54th edition
May. 2025	55th edition
Jul. 2025	56 th edition

Change notes

52nd edition

- Trademarks updated
- Add description of restricted access yum repository configuration file
- Add Confirm VE device path information file

53rd edition

- Support of Aurora Software for RHEL/CentOS 7.9 and RHEL/Rocky Linux 8.6 is ended
- RHEL/Rocky Linux 8.10 is supported
- MLNX_OFED 23.10-3.2.2.0 is supported
- Add description of configuration for persistent device name of HCAs
- Change procedure to update VMC Firmware
- Add procedure to update HCA Firmware
- Add description for installation packages by Vector Engine Type

54th edition

- Add description for model B403-8

55th edition

- Support of Aurora Software for RHEL/Rocky Linux 8.8 is ended

56th edition

- Minor wording adjustments

Index

SX-Aurora TSUBASA System Software

SX-Aurora TSUBASA

Installation Guide

56th Edition Jul. 2025

NEC Corporation

© 2018 NEC Corporation
No part of this document may be reproduced, in any form or by any means, without permission from NEC Corporation.