

IBM Spectrum NAS
Version 1.7.1.0

Installation Guide



Introduction

This manual will guide you through the installation of IBM Spectrum NAS firmware on each storage node. The installation is mostly automated and requires only a small amount of user input.

After completing the process, you will have the storage nodes prepared and ready for being added to an IBM Spectrum NAS storage cluster. Repeat the steps for each node that you want to use.

This edition applies to IBM Spectrum NAS, Version 1.7.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2018.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

System Requirements.....	4
Installation	6
Step 1: Insert the CD / ISO	6
Step 2: Reboot to start the installation	6
Step 3: Set a name	7
Step 4: Select system disk	7
Step 5: Select extras.....	8
Getting Help.....	10
IBM Spectrum NAS Support	10
Frequently Asked Questions (FAQ).....	11

System Requirements

Before you begin the installation of the IBM Spectrum NAS storage solution, please make sure that your storage nodes, minimum four, meet the following system requirements:

	<i>Recommended</i>	<i>Note</i>
CPU	x86-64 (4+ cores) with SSE4.1	One or two CPUs depending on use case
RAM Memory	64+ GB	Minimum 32 GB
Boot disk	60+ GB	Minimum 20 GB
Cache disk	100+ GB NVMe (one or more)	SAS/SATA is not supported for cache
Storage disk	SAS/SATA disks (multiple disks)	
Controller card	HBA (IT mode)	
Network	10/25/40 Gigabit NIC or better	Minimum 2x physical networks
Number of nodes	Scale-out from 4 nodes	The cluster can run with 3 nodes
Network switch	10/25/40 Gigabit switches	Multiple switches for redundancy

Notes:

Number of nodes	The minimum number of nodes is determined by the Resilience level and Erasure Coding to be used, equal to $n + 2$. For example: With Erasure Coding 2+2 (resilience 2), the minimum is four. The recommended minimum for Resilience 2 is five nodes, because a four-node cluster cannot continue to operate upon losing two nodes due to more than 50% of the gateways are being required.
OS	OS is not required for the nodes, it will be installed together with the IBM Spectrum NAS software.
IGMP	IGMP is required for all switches that handle the cluster's private and management networks. For details, see the IBM Spectrum NAS Quick Setup Guide or Management Manual.
Cache disks	For cache, NVMe devices are required, 100 GB or more, one or more devices each node. SAS/SATA SSD is not supported for cache due to lack of performance. Cache is shared for reads and writes and you can select, after setup, the amount used for each. Note: To avoid delays due to excessive cache misses (data not available in cache) thus moving data in and out of cache, the total cache size for all nodes combined should be larger than the size of your active data set. Consider overhead for (write) cache replication in order to secure updates. For durability, high endurance enterprise-grade NVMe (DWPD > 9) should be used.
Boot disk	The boot disk holds system log files in addition to OS and firmware. The recommended size is sufficient as no more than a few GB of truncated log files are usable. Mirrored boot disks (RAID1) is an option when availability is highly important, to lower the risk of having to reinstall a node and reinitialize the data disks. Please check information on RAID compatibility before implementation.
RAM	The amount of RAM recommended for each node depends on the amount of files per node, to keep most of metadata in RAM. Large nodes holding millions of files may benefit from up to or beyond 96 GB of RAM.

Storage disks	The number of storage disks will have a direct effect on the maximum sustainable I/O performance of the cluster. Example: twice as many storage disks will reduce the reads and writes to each disk by half and thus improve the total performance. Thus, many disks are better than a few very large disks. Make sure there is sufficient usable storage capacity available to allow the self-healing processes to recover data, should one or more nodes go down.
RAID	RAID should not be used in the IBM Spectrum NAS storage solution as it would introduce complexity and issues without advantages. With RAID, Identifying a defective disk within a node would be inconvenient or time consuming as the node would have to be booted into RAID BIOS or using tools not part of the solution. Instead, redundancy lies within the system through replication, copies and erasure coding. The recommendation is to use HBA for the storage devices, in IT mode only.
Resilience level	The resilience level is selected after installation, when creating a file system, and cannot be altered later. If you plan on using the cluster in production, resilience 2 should be used. In combination with the Virtual IP feature, this allows for safe non-interruptive rolling updates. For more details, see the Quick Setup Guide after installation, or the Management Manual.

**Note:**

If the disks you put in the node have been previously attached to a RAID card, and you do not intend to use RAID with the new installation, make sure that the disks are formatted. Otherwise, the old RAID metadata may interfere with the new installation (see the [FAQ section](#) for details).

Installation

Repeat the following steps for each node you want to install.

Step 1: Insert the CD / ISO

Turn on the node and insert the installation media in your source drive.

The ISO installation file has to be mounted to each node as a virtual drive, or transferred to bootable optical media, for example CD. USB flash memory and similar is not supported.

Typically, nodes are not set up to automatically boot from an external source drive. The methods for booting from external drive may vary, so you may have to refer to the documentation for your system's BIOS/UEFI.

Alternatively, try one of the following:

- Pressing F10 or F12 during initial system start-up, will normally allow you to enter a one-time boot menu. This will give you options to boot from external drive.
- Look out for an initial system startup message that indicates how to enter the BIOS/UEFI setup, often the Delete key. Alter the Boot Sequence from within the BIOS/UEFI.

Step 2: Reboot to start the installation

Boot the node from the source drive.

The installer will initialize and present the progress, similar to this:

```
[ OK ] Reached target Local File Systems (Pre).
[ OK ] Started udev Wait for Complete Device Initialization.
      Starting Activation of DM RAID sets...
[ OK ] Started Activation of DM RAID sets.
[ OK ] Reached target Encrypted Volumes.
[ OK ] Reached target Local File Systems.
      Starting Tell Plymouth To Write Out Runtime Data...
      Starting Import network configuration from initramfs...
      Starting Rebuild Journal Catalog...
[ OK ] Started Tell Plymouth To Write Out Runtime Data.
[ OK ] Started Rebuild Journal Catalog.
      Starting Update is Completed...
[ OK ] Started Update is Completed.
[ OK ] Started Import network configuration from initramfs.
      Starting Create Volatile Files and Directories...
[ OK ] Started Create Volatile Files and Directories.
      Starting Update UTMP about System Boot/Shutdown...
[ OK ] Started Update UTMP about System Boot/Shutdown.
[ OK ] Reached target System Initialization.
[ OK ] Listening on D-Bus System Message Bus Socket.
[ OK ] Listening on Open-iSCSI iscsiuid Socket.
[ OK ] Listening on Open-iSCSI iscsid Socket.
[ OK ] Reached target Sockets.
[ OK ] Reached target Basic System.
      Starting Dump dmesg to /var/log/dmesg...
      Starting Service enabling compressing RAM with zRam...
      Starting Login Service...
[ OK ] Started Hardware RNG Entropy Gatherer Daemon.
      Starting Hardware RNG Entropy Gatherer Daemon...
[ OK ] Started Anaconda NetworkManager configuration.
      Starting Anaconda NetworkManager configuration...
      Starting firewalld - dynamic firewall daemon...
      Starting Wait for Plymouth Boot Screen to Quit...
      Starting pre-anaconda logging service...
      Starting Terminate Plymouth Boot Screen...
[ OK ] Started Service enabling compressing RAM with zRam.
```

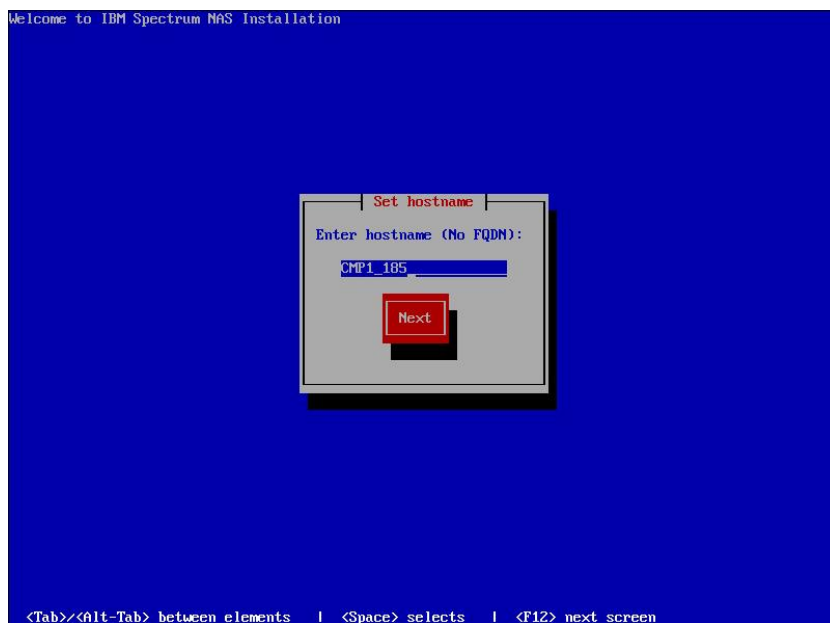


Note: If the disks you put in the node have been previously attached to a RAID card, and you do not intend to use RAID with the new installation, make sure that the disks are formatted. Otherwise, the old RAID metadata may interfere with the new installation (see the [FAQ section](#) for details).

Step 3: Set a name

Set the hostname for the node and press **Enter** twice to proceed. Hostname is a label assigned to the node that will later be used to identify it in the IBM Spectrum NAS Management Tool.


Note: No FQDN (fully qualified domain name) and no "dots" in the host name.



Warning: The installer will automatically format the disk after the next step when you select the system disk. This will permanently delete all data on the drive. Back up any files and settings that you want to keep, if you plan to restore them after the installation is complete.

Step 4: Select system disk


The installation wizard will list all available disks in the node. Select the disk that should be used as a boot / system disk for the node, and press **Enter**.



Step 5: Select extras

These options only apply if you are using Mellanox NIC's. If that is the case, then select (Space or Enter).

Tab down to the Finish button and press Enter to proceed.



The installer proceeds to install the packages of the IBM Spectrum NAS solution. At this point, no further user input is required until all packages have been installed. The time required for the installation depends on the number of packages and the speed of the hardware.

```

Welcome to IBM Spectrum NAS Installation

14:40:57 Not asking for UNC because of an automated install
14:40:57 Not asking for UNC because text mode was explicitly asked for in kickstart
14:40:57 Not asking for UNC because we don't have a network
Starting automated install.....
Checking software selection
Generating updated storage configuration
Checking storage configuration...
=====
Installation
1) [x] Language settings          2) [x] Time settings
   (English (United States))      (Etc/UTC timezone)
3) [x] Installation source        4) [x] Software selection
   (Local media)                  (Custom software selected)
5) [x] Installation Destination   6) [x] Kdump
   (Custom partitioning selected) (Kdump is disabled)
7) [ ] Network configuration      8) [ ] User creation
   (Not connected)                (No user will be created)
=====
Progress
Setting up the installation environment
.
Creating disklabel on /dev/sda
.
Creating ext4 on /dev/sda2

```



Note: If the screen looks different than above, refer to the FAQ section:
[My node installation doesn't look like described, what could be wrong.](#)

After the installation is completed, the node will automatically reboot and bring you to the login screen. There is no need to log in, you can leave the node as it is. A second reboot will occur after a few minutes, and only then the node is ready for use.

```
CentOS Linux 7 (Core)
Kernel 4.9.76 on an x86_64

CMP1_185 login: _
```

Repeat the procedure for all nodes you want to install: [Step 1](#)



Note: If the node has booted from the installation media, instead of the system disk – i.e. you see the screen above – do the following:

1. Remove any installation media, if not automatically ejected upon reboot.
 2. Make sure that the default boot drive is the system disk. You may want to change the boot sequence in BIOS/UEFI.
-

Getting Help

IBM Spectrum NAS Support

IBM Spectrum NAS's support service is always available to answer your questions by email, phone or web, based on your support offering.

To get help with troubleshooting, please contact your technical service representative.

Your location	Method of contacting the IBM Support Center
In the United States	Call 1-800-IBM-SERV for support.
Outside the United States	Contact your local IBM Support Center or see the Directory of worldwide contacts .

Please note that providing us with background information will help us to do preliminary research to understand your issue better, to make a more efficient interaction. We would therefore ask you to provide the following information when you send a support request:

- Cluster ID
- Software version of all relevant software
- An approximation of time since the issue first occurred
- Whether the issue is reproducible
- Steps taken so far to solve the problem

Frequently Asked Questions (FAQ)

The node installation CD/DVD does not start

Reboot the system into BIOS and verify if the booting sequence has been changed to the installation media drive. Save configuration and continue installation.

My node installation doesn't look like described, what could be wrong?

Alternative 1: If you, at Step 5, see this screen:

```
09:57:38 Not asking for UMC because of an automated install
09:57:38 Not asking for UMC because text mode was explicitly asked for in kickstart
09:57:38 Not asking for UMC because we don't have a network
Starting automated install...
Generating updated storage configuration
storage configuration failed: The following problem occurred on line 2 of the kickstart file:
Requested boot drive sdh doesn't exist or cannot be used
=====
Installation
1) [x] Timezone settings                2) [x] Installation source
   (Etc/UTC timezone)                  (Local media)
3) [x] Software selection                4) [!] Install Destination
   (Custom software selected)          (No disks selected)
5) [x] Network settings                 6) [x] Create user
   (Wired (enp4s0f0) connected)        (No user will be created)
Please make your choice from above [ 'q' to quit | 'b' to begin installation |
'r' to refresh]:
```

The boot disk that is connected to the SATA interface contains old RAID metadata. This usually happens when a disk was previously configured in a RAID setup and then connected to SATA without “zeroing out” the RAID metadata. Linux tries to identify the disk but the installation does not support this.

Our way of clearing the RAID metadata is to connect the disk to a windows PC, then use a tool e.g. Western Digital's "Data Lifeguard Diagnostic for Windows" and perform a quick wipe "write zeros" to the disk.

This tool can be found at Western Digital download site: <http://support.wdc.com/downloads.aspx>

Alternative 2: If you, at Step 5, see this screen:

```
payload.preinstall(packages=packages, groups=payload.languageGroups())
File "/usr/lib64/python2.7/threading.py", line 764, in run
self._target(*self._args, **self._kwargs)
File "/usr/lib64/python2.7/site-packages/pyanaconda/threads.py", line 211, in run
threading.Thread.run(self, *args, **kwargs)
NoSuchPackage: mdadm
What do you want to do now?
1) Report Bug
2) Debug
3) Quit
Please make your choice from above:
```

You find this screen when you have a problem similar to the one above. There is no problem with the boot disk, the problem is caused by another disk in the node that has been attached to a RAID card earlier and therefore contains RAID metadata. Find the disk containing the RAID metadata and clear it using the same procedure as described above, Case 1.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM® product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp.

Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of the Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to

collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.



Printed in USA

GC27-9226-01

