



Setup wizard

SANtricity 11.5

NetApp
October 22, 2024

Table of Contents

- Setup wizard 1
- Concepts 1
- FAQs 2

Setup wizard

Concepts

Setup wizard overview

Use the Setup wizard to configure your storage array, including hardware, hosts, applications, workloads, pools, alerts, and AutoSupport.

When you open System Manager for the first time, the Setup wizard launches.

The Setup wizard prompts you to perform basic configuration tasks, such as naming your storage array, configuring your hosts, selecting applications, and creating pools of storage.

If you cancel the wizard, you cannot manually relaunch it.

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Setup wizard terminology

Learn how the Setup wizard terms apply to your storage array.

Term	Description
Application	An application is a software program, such as Microsoft SQL Server or Microsoft Exchange.
Alert	Alerts notify administrators about important events that occur on the storage arrays. Alerts can be sent via email, SNMP traps, or syslog.
AutoSupport	The AutoSupport feature monitors the health of a storage array and sends automatic dispatches to technical support.
Hardware	The storage system hardware includes storage arrays, controllers, and drives.
Host	A host is a server that sends I/O to a volume on a storage array.

Term	Description
Object	An object is any logical or physical storage component. Logical objects include volume groups, pools, and volumes. Physical objects include the storage array, array controllers, hosts, and drives.
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Volume	<p>A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.</p> <p>A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.</p>
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

FAQs

What if I don't see all of my hardware components?

If you do not see all your hardware components on the **Verify hardware** dialog box, it could mean that a drive shelf is not connected correctly, or that an incompatible shelf is installed in the storage array.

Verify that all drive shelves are connected correctly. If you are uncertain about which drive shelves are compatible, contact technical support.

What if I don't see all of my hosts?

If you do not see your connected hosts, then automatic detection has failed, the hosts are improperly connected, or no hosts are currently connected.

You can configure hosts later, once you are done with the setup. You can create hosts either automatically or manually as follows:

- If you installed the Host Context Agent (HCA) on your hosts, the HCA pushes the host configuration information to the storage array. System Manager automatically configures these hosts and displays them in the **Initial Setup** wizard.
- You can manually create hosts and associate the appropriate host port identifiers by going to **Storage > Hosts**. Hosts that have been created manually also display in the **Initial Setup** wizard.
- The target and host must be configured for the host port type (for example, iSCSI or NVMe over InfiniBand), and a session to the storage established before automatic detection will work.

How do I know which host operating system type is correct?

The Host Operating System Type field contains the operating system of the host. You can select the recommended host type from the drop-down list or allow the Host Context Agent (HCA) to configure the host and appropriate host operating system type.

Host Operating System type	Operating System (OS) and multipath driver
AIX MPIO	The Advanced Interactive Executive (AIX) OS and the native MPIO driver
AVT_4M	Silicon Graphics, Inc. (SGI) proprietary multipath driver; refer to the SGI installation documentation for more information
Factory Default	This is reserved for the initial startup of the storage array and should be changed to match the host operating system and multipath driver being used for the particular host
HP-UX	The HP-UX OS with native multipath driver
Linux (ATTO)	The Linux OS and the ATTO Technology, Inc. driver (must use ATTO FC HBAs)
Linux (DM-MP)	The Linux OS and the native DM-MP driver
Linux (Pathmanager)	The Linux OS and the SGI proprietary multipath driver; refer to the SGI installation documentation for more information

Host Operating System type	Operating System (OS) and multipath driver
Mac OS	The Mac OS and the ATTO Technology, Inc. driver
ONTAP	FlexArray
Solaris (version 11 or later)	The Solaris 11 or later OS and the native MPxIO driver
Solaris (version 10 or earlier)	The Solaris 10 or earlier OS and the native MPxIO driver
SVC	IBM SAN Volume Controller
VMware	The ESXi OS
Windows or Windows Clustered	The Windows Server OS and Windows MPIO with a DSM driver
Windows (ATTO)	The Windows OS and the ATTO Technology, Inc. driver

After the HCA is installed and the storage is attached to the host, the HCA sends the host topology to the storage controllers through the I/O path. Based on the host topology, the storage controllers automatically define the host and the associated host ports, and then set the host type.



If the HCA does not select the recommended host type, you must manually set the host type in System Manager.

How does identifying applications help me manage my storage array?

When you identify applications, System Manager automatically recommends a volume configuration that optimizes storage based on application type.

Optimizing volumes by application can make data storage operations more efficient. Characteristics such as I/O type, segment size, controller ownership, and read and write cache are included in the volume configuration. In addition, you can view performance data by application and by workload to assess the latency, IOPS, and MiB/s of applications and their associated workloads.

What is a workload?

For some applications in your network, such as SQL Server or Exchange, you can define a workload that optimizes storage for that application.

A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume

characteristics are optimized to support Microsoft SQL Server.

During volume creation, System Manager prompts you to answer questions about a workload's use. For example, if you are creating volumes for Microsoft Exchange, you are asked how many mailboxes you need, what your average mailbox capacity requirements are, and how many copies of the database you want. System Manager uses this information to create an optimal volume configuration for you, which can be edited as needed.

How do I configure SNMP or syslog alerts?

In addition to email alerts, you can configure alerts to be sent by Simple Network Management Protocol (SNMP) traps or by syslog messages.

To configure SNMP or syslog alerts, go to **Settings > Alerts**.

How do I configure the delivery method for AutoSupport?

To access configuration tasks for AutoSupport delivery methods, go to **Support > Support Center**, and then click the AutoSupport tab.

The following protocols are supported: HTTPS, HTTP, and SMTP.

What type of data is collected through the AutoSupport feature?

The AutoSupport feature contains three standard dispatch types: event dispatches, scheduled dispatches, and on-demand and remote diagnostics dispatches.

The AutoSupport data does not contain any user data.

- **Event dispatches**

When events occur on the system that warrant proactive notification to technical support, the AutoSupport feature automatically sends an event-triggered dispatch.

- Sent when a support event on the managed storage array occurs.
- Includes a comprehensive snapshot of what was going on with the storage array at the time the event occurred.

- **Scheduled dispatches**

The AutoSupport feature automatically sends several dispatches on a regular schedule.

- **Daily dispatches** — Sent once every day during a user-configurable time interval. Includes the current system event logs and performance data.
- **Weekly dispatches** — Sent once every week during a user-configurable time interval and day. Includes configuration and system state information.

- **AutoSupport OnDemand and Remote Diagnostics dispatches**

- **AutoSupport OnDemand** — Allows technical support to request retransmission of a previous AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending retransmission requests and responds accordingly.

- **Remote Diagnostics** — Allows technical support to request a new, up-to-date AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending new requests and responds accordingly.

How do I know if I should accept the recommended pool configuration?

Whether you accept the recommended pool configuration depends on a few factors.

Determine the type of storage that is best for your requirements by answering the following questions:

- Do you prefer multiple pools of smaller capacities, rather than the largest pools possible?
- Do you prefer RAID volume groups over pools?
- Do you prefer to manually provision your drives, rather than having a configuration recommended for you?

If you answered Yes to any of these questions, consider rejecting the recommended pool configuration.

System Manager has not detected any hosts. What do I do?

If you do not see your connected hosts, then automatic detection has failed, the hosts are improperly connected, or no hosts are currently connected.

You can configure hosts later, once you are done with the setup. You can create hosts either automatically or manually as follows:

- If you installed the Host Context Agent (HCA) on your hosts, the HCA pushes the host configuration information to the storage array. System Manager automatically configures these hosts and displays them in the **Initial Setup** wizard.
- You can manually create hosts and associate the appropriate host port identifiers by going to **Storage > Hosts**. Hosts that have been created manually also display in the **Initial Setup** wizard.
- The target and host must be configured for the host port type (for example, iSCSI or NVMe over InfiniBand), and a session to the storage established before automatic detection will work.

What is a hot spare drive?

Hot spares act as standby drives in RAID 1, RAID 5, or RAID 6 volume groups. They are fully functional drives that contain no data. If a drive fails in the volume group, the controller automatically reconstructs data from the failed drive to a hot spare.

If a drive fails in the storage array, the hot spare drive is automatically substituted for the failed drive without requiring a physical swap. If the hot spare drive is available when a drive fails, the controller uses redundancy data to reconstruct the data from the failed drive to the hot spare drive.

A hot spare drive is not dedicated to a specific volume group. Instead, you can use a hot spare drive for any failed drive in the storage array with the same capacity or smaller capacity. A hot spare drive must be of the same media type (HDD or SSD) as the drives that it is protecting.



Hot spare drives are not supported with pools. Instead of hot spare drives, pools use the preservation capacity within each drive that comprises the pool.

What is a volume group?

A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)

What are the differences between pools and volume groups?

A pool is similar to a volume group, with the following differences.

- The data in a pool is stored randomly on all drives in the pool, unlike data in a volume group, which is stored on the same set of drives.
- A pool has less performance degradation when a drive fails, and takes less time to reconstruct.
- A pool has built-in preservation capacity; therefore, it does not require dedicated hot spare drives.
- A pool allows a large number of drives to be grouped.
- A pool does not need a specified RAID level.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.