



Complete the expansion

System Manager Classic

NetApp
June 22, 2024

Table of Contents

- Complete the expansion. 1
 - Configure the node details in System Manager 1
 - Configure AutoSupport on the new nodes 1
 - Configure the Service Processor network 2
 - Validate the configuration of the expanded cluster 3
 - Generate an AutoSupport message about completing expansion 5
 - Update LUN paths for the new nodes 6

Complete the expansion

After both nodes are joined to the cluster, you must finish configuring the newly added nodes by configuring AutoSupport and completing the SP network. You then validate the expanded cluster and generate an AutoSupport message to complete the expansion. If the cluster uses SAN, you should update LUN paths.

Configure the node details in System Manager

You can use System Manager to configure the node management LIF and Service Processor settings for the newly added nodes.

Before you begin

- Sufficient number of ports must be present in the default IPspace for LIF creation.
- All the ports must be up and running.

Steps

1. Configure node management:
 - a. Enter the IP address in the **IP Address** field.
 - b. Select the port for node management in the **Port** field.
 - c. Enter the netmask and gateway details.
2. Configure Service Processor settings:
 - a. Select the **Override defaults** check box to override the default values.
 - b. Enter the IP address, netmask, and gateway details.
3. Click **Submit and Proceed** to complete the network configuration of the nodes.
4. Verify the details of the nodes in the **Summary** page.

What to do next

- If your cluster is protected, you should create the required number of intercluster LIFs in the newly added nodes to avoid partial peering and unhealthy protection.
- If SAN data protocols are enabled in your cluster, you should create the required number of SAN Data LIFs for serving data.

Configure AutoSupport on the new nodes

After you add nodes to a cluster, you must configure AutoSupport on the nodes.

Before you begin

AutoSupport must be set up on the cluster's existing nodes.

About this task

You must perform this procedure on both the nodes.

Steps

1. View the AutoSupport configuration using the `system node autosupport show` command with the `-node` parameter set to one of the nodes in the original cluster.

```
cluster1::> system node autosupport show -node cluster1-1
                Node: cluster1-1
                State: enable
                SMTP Mail Hosts: smtp.example.com
...

```

2. On one of the newly added nodes, configure AutoSupport in the same way that it is configured on the existing nodes by using the `system node autosupport modify` command.

```
cluster1::> system node autosupport modify -node cluster1-3 -state
enable -mail-hosts smtp.example.com -from alerts@node3.example.com -to
support@example.com -support enable -transport https -noteto
pda@example.com -retry-interval 23m

```

3. Repeat the previous step for the other newly added node.

Configure the Service Processor network

After you expand a cluster, you must configure the Service Processor (SP) network on the new nodes. If the SP uses manual network configuration, you must configure the IP addresses for the SP on the new nodes. If the SP uses automatic network configuration, you must identify the IP addresses that were selected.

Steps

1. If the cluster SP uses manual network configuration, configure IP addresses on both nodes for the SP network by using the `system service-processor network modify` command.

The following commands configure the SP network in cluster1-3 and cluster1-4 nodes:

```
cluster1::> system service-processor network modify -node cluster1-3
-address-family IPv4 -enable true -ip-address 192.168.123.98-netmask
255.255.255.0 -gateway 192.168.123.1
cluster1::> system service-processor network modify -node cluster1-4
-address-family IPv4 -enable true -ip-address 192.168.123.99 -netmask
255.255.255.0 -gateway 192.168.123.1

```

2. Verify that the SP network is configured correctly on both the new nodes by using the `system service-processor network show` command for each node.

The status should be succeeded. Verification is required in all situations. Even if the SP network was

automatically configured, you should verify that it was configured successfully, and you must determine which IP addresses were assigned.

The following output indicates that both the cluster1-3 and the cluster1-4 nodes have successful SP network setup:

```
cluster1::> system service-processor network show -node cluster1-3
Address
Node          Status          Family    Link State  IP Address
-----
cluster1-3    online          IPv4      up           192.168.123.98

                DHCP: none
                MAC Address: 00:a0:98:43:a1:1e
                Network Gateway: 10.60.172.1
                Network Mask (IPv4 only): 255.255.255.0
                Prefix Length (IPv6 only): -
                IPv6 RA Enabled: -
                Subnet Name: -
                SP Network Setup Status: succeeded
                ...

cluster1::> system service-processor network show -node cluster1-4
Address
Node          Status          Family    Link State  IP Address
-----
cluster1-4    online          IPv4      up           192.168.123.99

                DHCP: none
                MAC Address: 00:a0:98:43:a1:1e
                Network Gateway: 10.60.172.1
                Network Mask (IPv4 only): 255.255.255.0
                Prefix Length (IPv6 only): -
                IPv6 RA Enabled: -
                Subnet Name: -
                SP Network Setup Status: succeeded
                ...
```

3. If your site typically has DNS entries for the SP network, verify that the DNS entries are created for the new nodes.

Validate the configuration of the expanded cluster

After you expand the cluster, you must validate the configuration by running Config

Advisor and using some commands that verify cluster health and cluster replication rings.

Steps

1. Check the health of the configuration by running Config Advisor:

- a. Start Config Advisor, and then click **Collect Data**.

Config Advisor displays any problems found.

- b. If problems are found, correct them and run the tool again.

2. Ensure that all nodes in the cluster are in a healthy state by using the `cluster show` command.

```
cluster-1::> cluster show
Node                Health  Eligibility
-----
cluster1-1          true    true
cluster1-2          true    true
cluster1-3          true    true
cluster1-4          true    true
4 entries were displayed.
```

3. Ensure that the cluster replication rings have the same epoch, database epoch, and database transaction numbers on all nodes in the cluster:

The easiest way to compare transaction numbers is to view them for one unit name at a time.

- a. Set the privilege level to advanced by using the `set -privilege advanced` command.
- b. View cluster ring information about the first unit name by using the `cluster ring show` command with the `-unitname mgmt` parameter, and verify that all nodes have the same number in the Epoch, DB Epoch, and DB Trnxs columns.

```

cluster-1::*> cluster ring show -unitname mgmt
Node          UnitName Epoch      DB Epoch DB Trnxs Master      Online
-----
cluster1-1
      mgmt      2          2          959      cluster1-1
                                           master
cluster1-2
      mgmt      2          2          959      cluster1-2
                                           secondary
cluster1-3
      mgmt      2          2          959      cluster1-3
                                           master
cluster1-4
      mgmt      2          2          959      cluster1-3
                                           secondary

4 entries were displayed.

```

- c. Repeat the command with the `-unitname vldb` parameter.
- d. Repeat the command with the `-unitname vifmgr` parameter.
- e. Repeat the command with the `-unitname bcomd` parameter.
- f. Repeat the command with the `-unitname crs` parameter.
- g. Return the privilege level to admin by using the `set -privilege admin` command.

Generate an AutoSupport message about completing expansion

After you expand a cluster, you should send an AutoSupport message to indicate that the expansion process is complete. This message communicates to internal and external support staff that the expansion is complete and acts as a timestamp for any troubleshooting that might be required later.

Before you begin

AutoSupport must be set up.

Steps

1. For each node in the cluster, send an AutoSupport message by using the `system node autosupport invoke` command.

You must issue the message once for each node in the cluster, including the newly added nodes.

If you added two nodes to a two-node cluster, you must send the message four times.

```
cluster1::> system node autosupport invoke -node * -message "cluster
expansion complete" -type all
The AutoSupport was successfully invoked on node "cluster1-1". To view
the status
of the AutoSupport, use the "system node autosupport history show"
command.
Note: It may take several minutes for the AutoSupport to appear in the
history list.
The AutoSupport was successfully invoked on node "cluster1-2". To view
the status
of the AutoSupport, use the "system node autosupport history show"
command.
Note: It may take several minutes for the AutoSupport to appear in the
history list.
The AutoSupport was successfully invoked on node "cluster1-3". To view
the status
of the AutoSupport, use the "system node autosupport history show"
command.
Note: It may take several minutes for the AutoSupport to appear in the
history list.
The AutoSupport was successfully invoked on node "cluster1-4". To view
the status
of the AutoSupport, use the "system node autosupport history show"
command.
Note: It may take several minutes for the AutoSupport to appear in the
history list.
4 entries were acted on.
```

Update LUN paths for the new nodes

If your cluster is configured for SAN, you must create SAN LIFs on the newly added nodes and then update paths.

About this task

This procedure is required only if the cluster contains LUNs. If the cluster contains only files, you can skip this procedure.

Steps

1. For each storage virtual machine (SVM) in the cluster, create new LIFs on the newly added nodes:
 - a. Identify the SVMs that use FC or iSCSI protocols by using the `vserver show` command with the `-fields allowed-protocols` parameter and reviewing the output.


```
cluster1::> vserver show -fields allowed-protocols
vserver allowed-protocols
-----
vs1      cifs,ndmp
vs2      fcp
vs3      iscsi
...
```

- b. For each SVM that uses FC or iSCSI, create at least two data LIFs on each of the newly added nodes by using the `network interface create` command with the `-role data` parameter.

```
cluster1::> network interface create -vserver vs1 -lif lif5 -role
data
-data-protocol iscsi -home-node cluster1-3 -home-port e0b
-address 192.168.2.72 -netmask 255.255.255.0
```

- c. For each SVM, verify that it has LIFs on all nodes in the cluster by using the `network interface show` command with the `-vserver` parameter.

2. Update portsets:

- a. Determine whether portsets exist by using the `lun portset show` command.
- b. If you want to make the new LIFs visible to existing hosts, add each new LIF to the portsets by using the `lun portset add` command—once for each LIF.

3. If you use FC or FCoE, update zoning:

- a. Verify that zoning is set up correctly to enable the existing initiator ports on the host to connect to the new target ports on the new nodes.
- b. Update switch zoning to connect the new nodes to existing initiators.

Zoning setup varies depending on the switch that you use.

- c. If you plan to move LUNs to the new nodes, expose the new paths to the hosts by using the `lun mapping add-reporting-nodes` command.

4. On all host operating systems, rescan to discover the newly added paths.
5. Depending on the host operating systems, remove any stale paths.
6. Add or remove paths to your MPIO configuration.

Related information

[SAN configuration](#)

[SAN administration](#)

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.