DM Series cabling and power on/off procedures

Cabling and power on/off procedures for the DM Series storage systems and disk shelves

Lenovo

The importance of proper storage system cabling

It is important to understand the cabling process and best practices. Many system issues and unplanned downtime can be attributed to poor or inaccurate cabling. Therefore, it is important to always verify the cabling before troubleshooting or making changes to a storage system.

Note: It is necessary to contact the network administrator for information about connecting the system to the switches.

Cluster networks

Depending on the environment, ONTAP clusters require one or more networks:

- Cluster interconnect: The cluster interconnect is a dedicated, redundant 10 GbE infrastructure connected between two nodes or connected to approved switches. In multinode clusters, nodes communicate with each other over a cluster interconnect. The interconnect enables the controllers to function as a single cluster. In a two-node cluster, the interconnect can be switchless. When more than two nodes are added to a cluster, a private cluster interconnect that uses switches is required.
- Management network: The management network is used for cluster administration. There are redundant connections to the management ports on each node and management ports on each cluster. An independent Ethernet switch should be utilized for the management network.
- Data network: The data network is used for clients and hosts to access data. The data network can be composed of one or more networks that are primarily used for data access by clients or hosts. Depending on the environment, there might be an Ethernet, an FC network, or a converged network. Networks can consist of one or more switches, or even redundant networks. In smaller environments, the management and data networks might be on a shared Ethernet network.

Cluster networks diagram

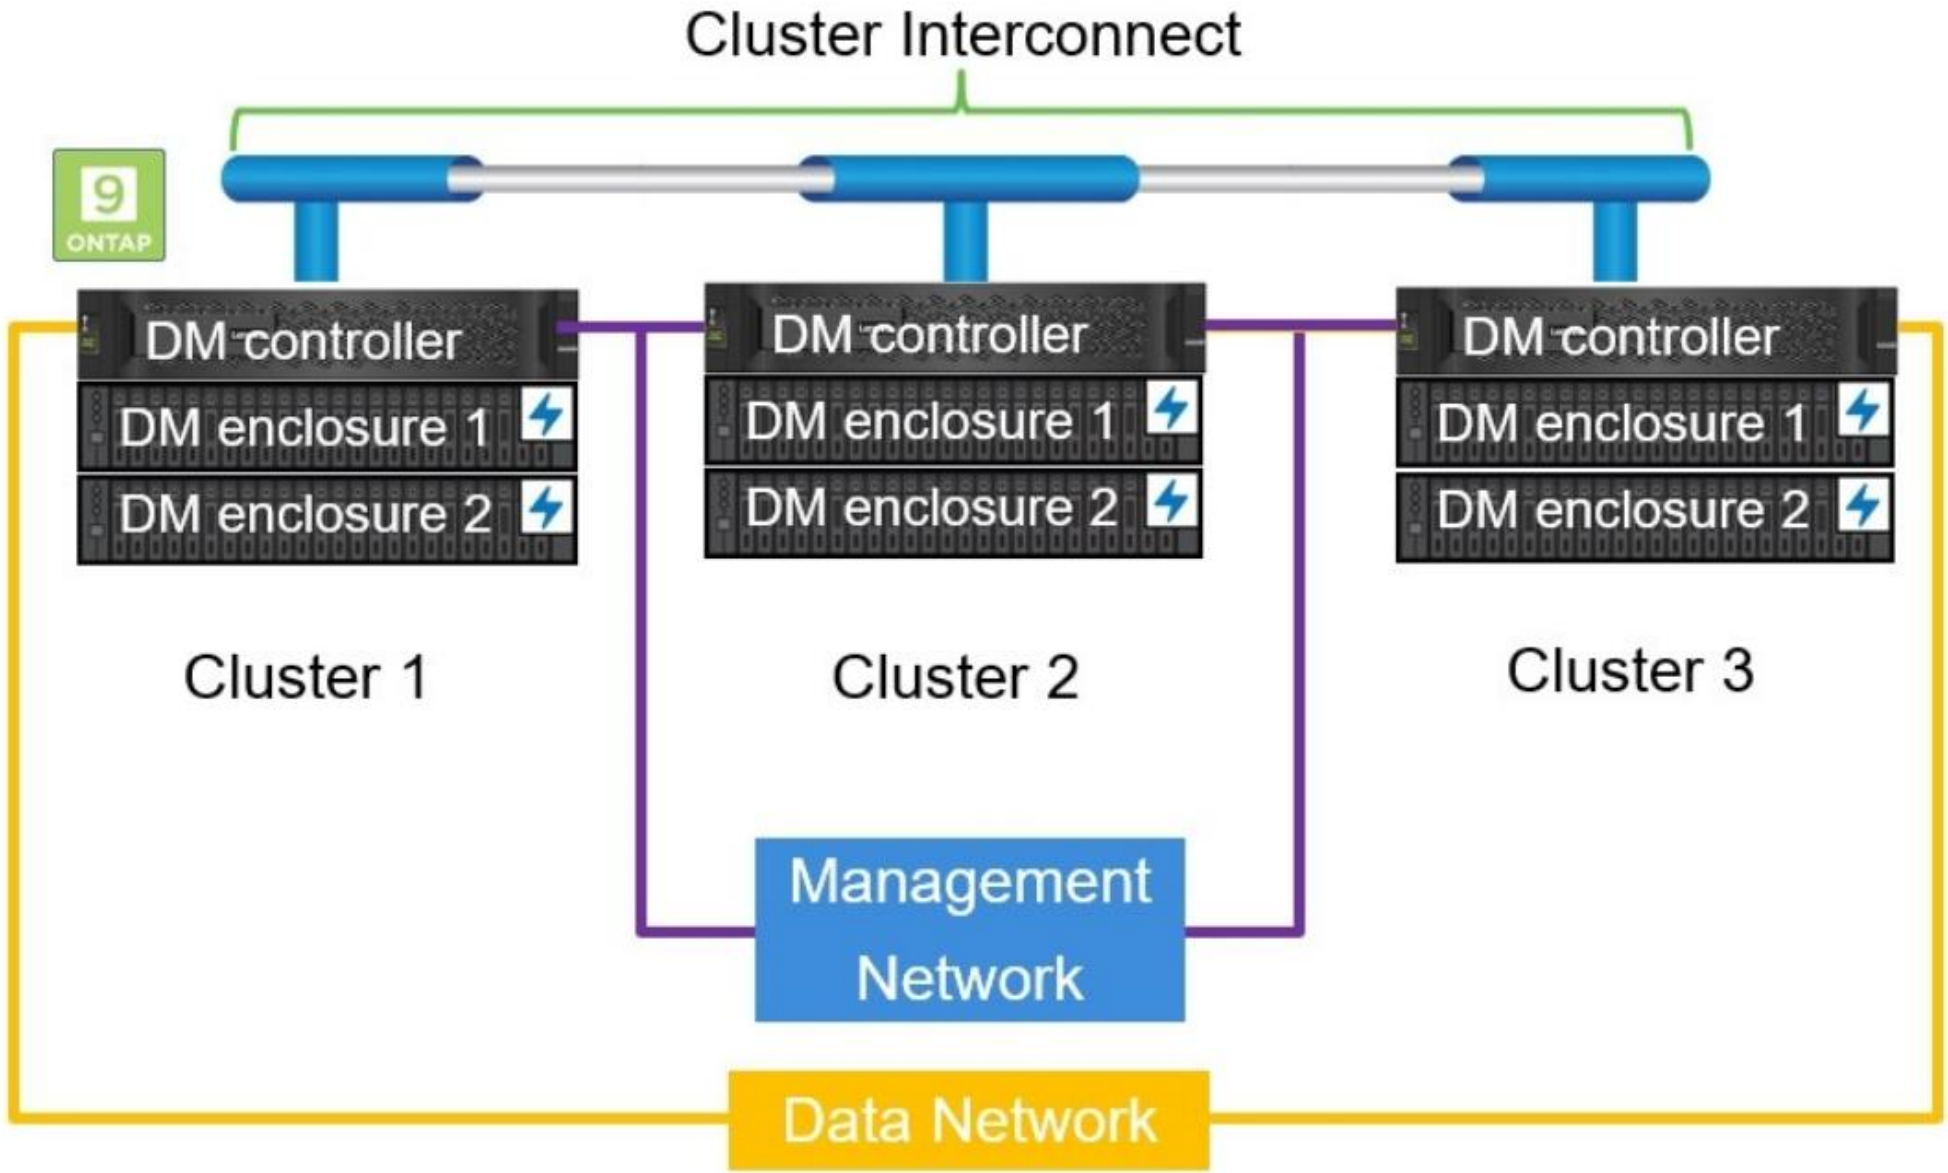

Cabling a two-node switchless cluster – DM3000/5000 UTA2 configuration

The management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers. Do not plug power cords in at this point.

1. Cable the cluster interconnect ports. Connect controller 1 e0a to controller 2 e0a and controller 1 e0b to controller 2 e0b.
2. Cable the UTA2 data ports e0c and e0d or e0e and e0f to the data network switches.
3. Cable the management ports to the management network switches.

Note: As you insert the connector, you should feel it clicks into place. If it does not click, remove it and try again.

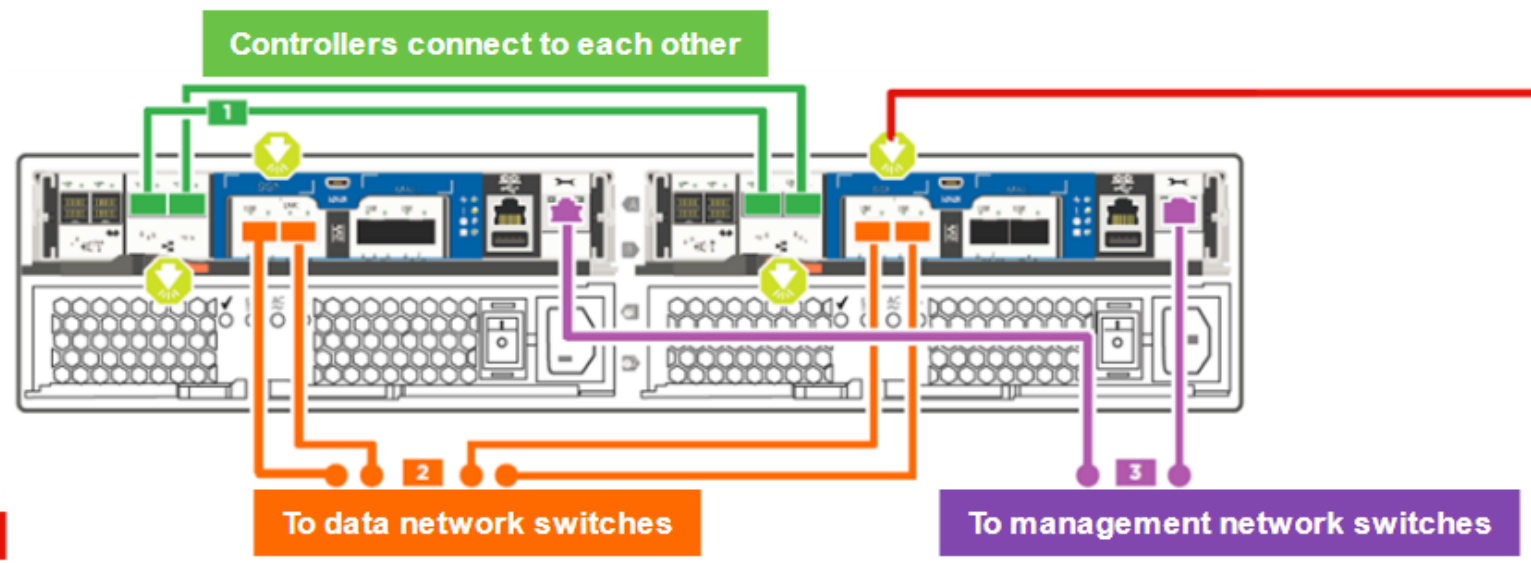

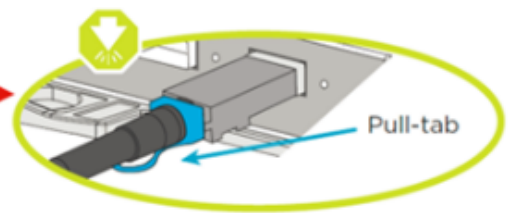

Be sure to check the 10 GbE network cable arrow for the proper cable connector pull-tab orientation before you begin connection.

Lenovo

Cabling a two-node switchless cluster – DM3000/5000 Ethernet configuration

The management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers. Do not plug power cords in at this point.

1. Cable the cluster interconnect ports. Connect controller 1 e0a to controller 2 e0a and controller 1 e0b to controller 2 e0b.
2. Cable the Ethernet data ports e0c and e0e or e0d and e0f to the data network switches.
3. Cable the management ports to the management network switches.

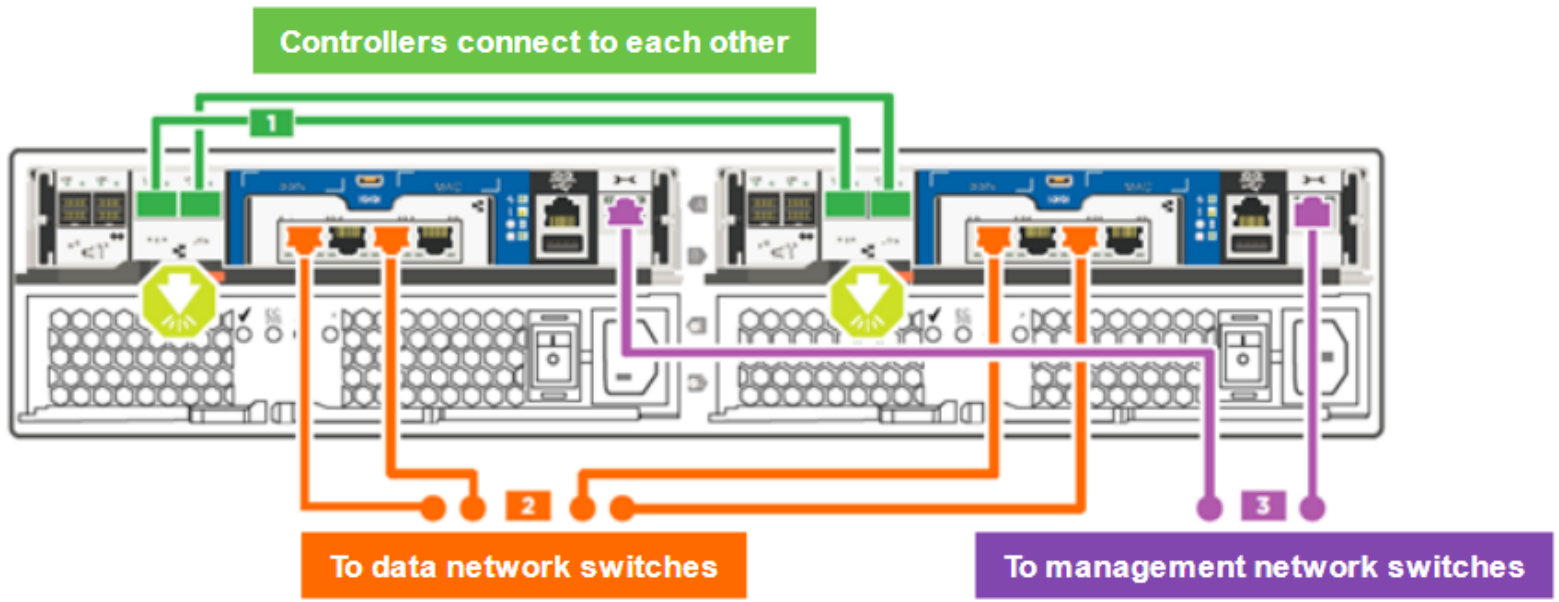

Cabling a two-node switchless cluster – DM7000 series

The management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

1. Cable the cluster interconnect ports. Connect controller 1 e0a to controller 2 e0a and controller 1 e0b to controller 2 e0b.
2. Cable the management ports to the management network switches.
3. Cable the Ethernet data ports e0e|0e through e0h|oh to the data network switches.
4. Cable the 10 GBASE-T ports e0c and e0d to the data network switches - for NAS or iSCSI connectivity.
5. Connect all the power cables (do not power on the system at this point).

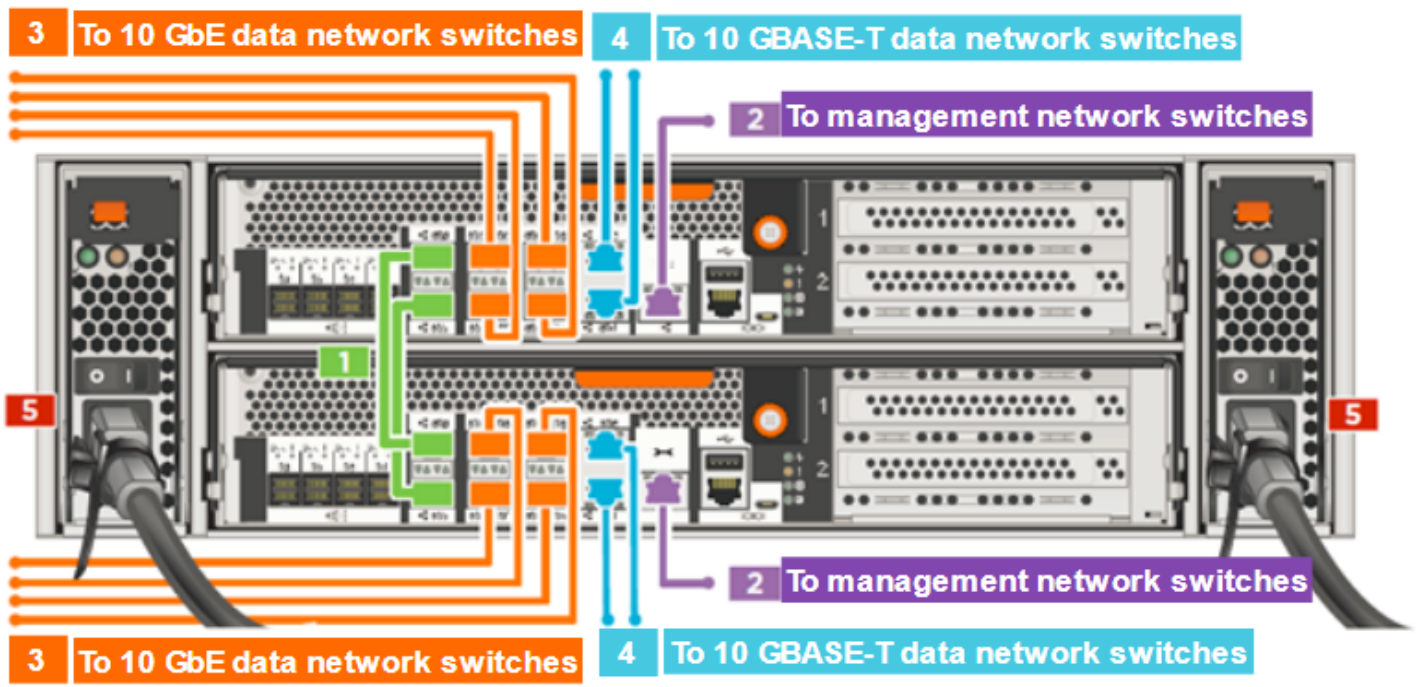

Cabling a switched cluster – DM3000/5000 UTA2 configuration

The management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches. Do not plug power cords in at this point.

1. Cable the cluster interconnect ports e0a and e0b of each controller to the cluster interconnect switches.
2. Cable the UTA2 data ports e0c and e0d or e0e and e0f to the data network switches.
3. Cable the management ports to the management network switches.

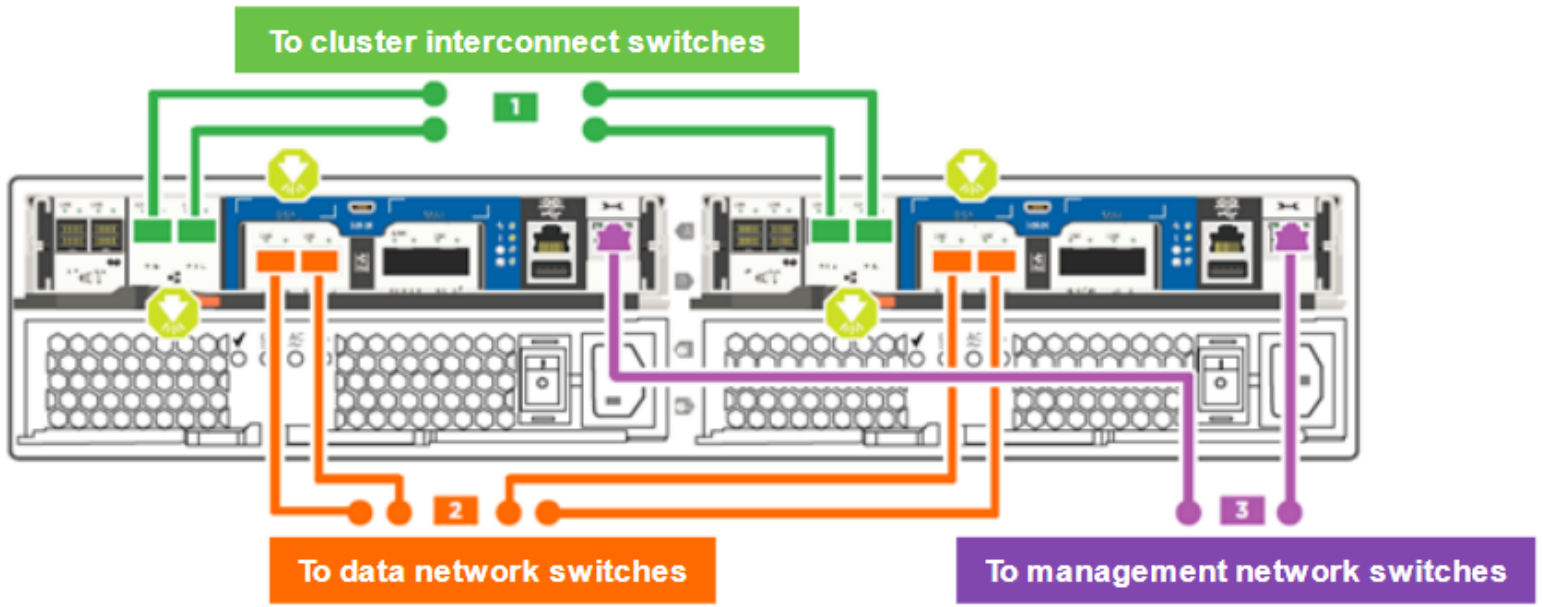

Cabling a switched cluster – DM3000/5000 Ethernet configuration

The management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches. Do not plug power cords in at this point.

1. Cable the cluster interconnect ports e0a and e0b of each controller to the cluster interconnect switches.
2. Cable the Ethernet data ports e0c and e0e or e0d and e0f to the data network switches.
3. Cable the management ports to the management network switches.

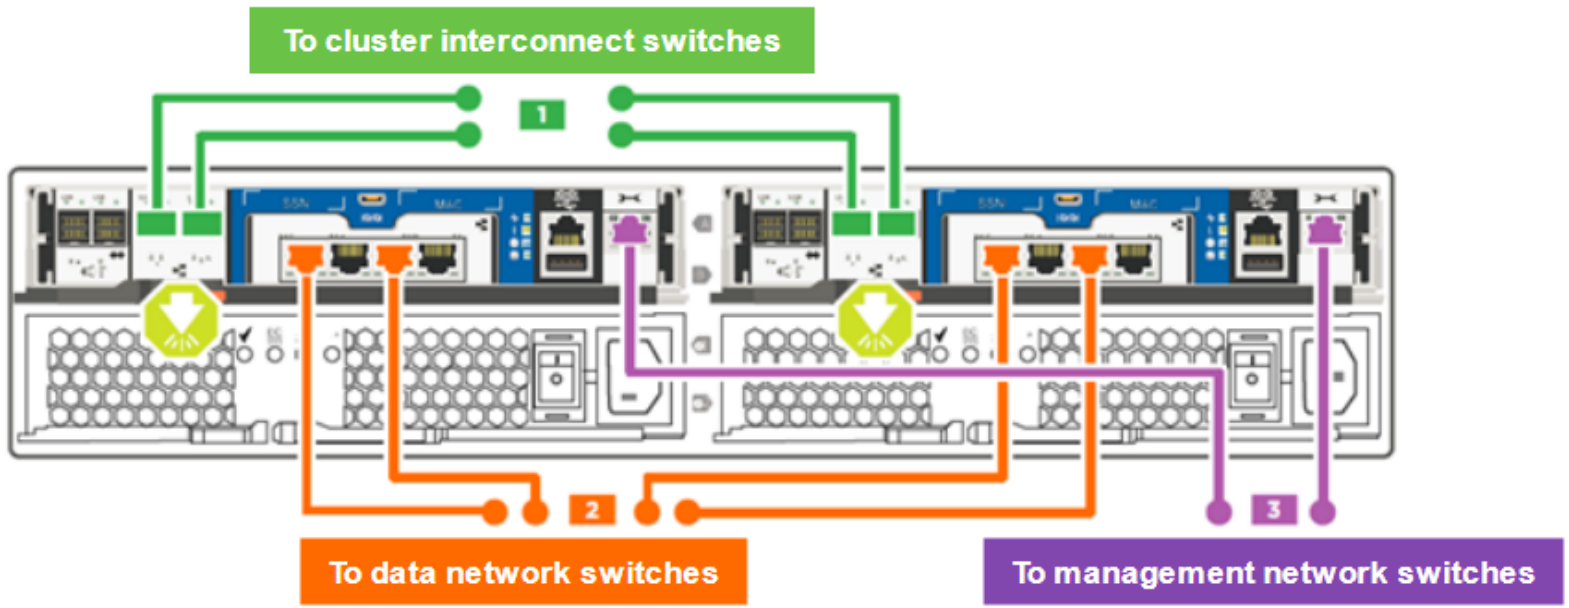

Cabling a two-node switched cluster – DM7000 series

The management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

1. Cable the cluster interconnect ports e0a and e0b of each controller to the cluster interconnect switches.
2. Cable the management ports to the management network switches.
3. Cable the Ethernet data ports e0e|0e through e0h|oh to the data network switches.
4. Cable the 10 GBASE-T ports e0c and e0d to the data network switches - for NAS or iSCSI connectivity.
5. Connect all the power cables (do not power on system at this point).

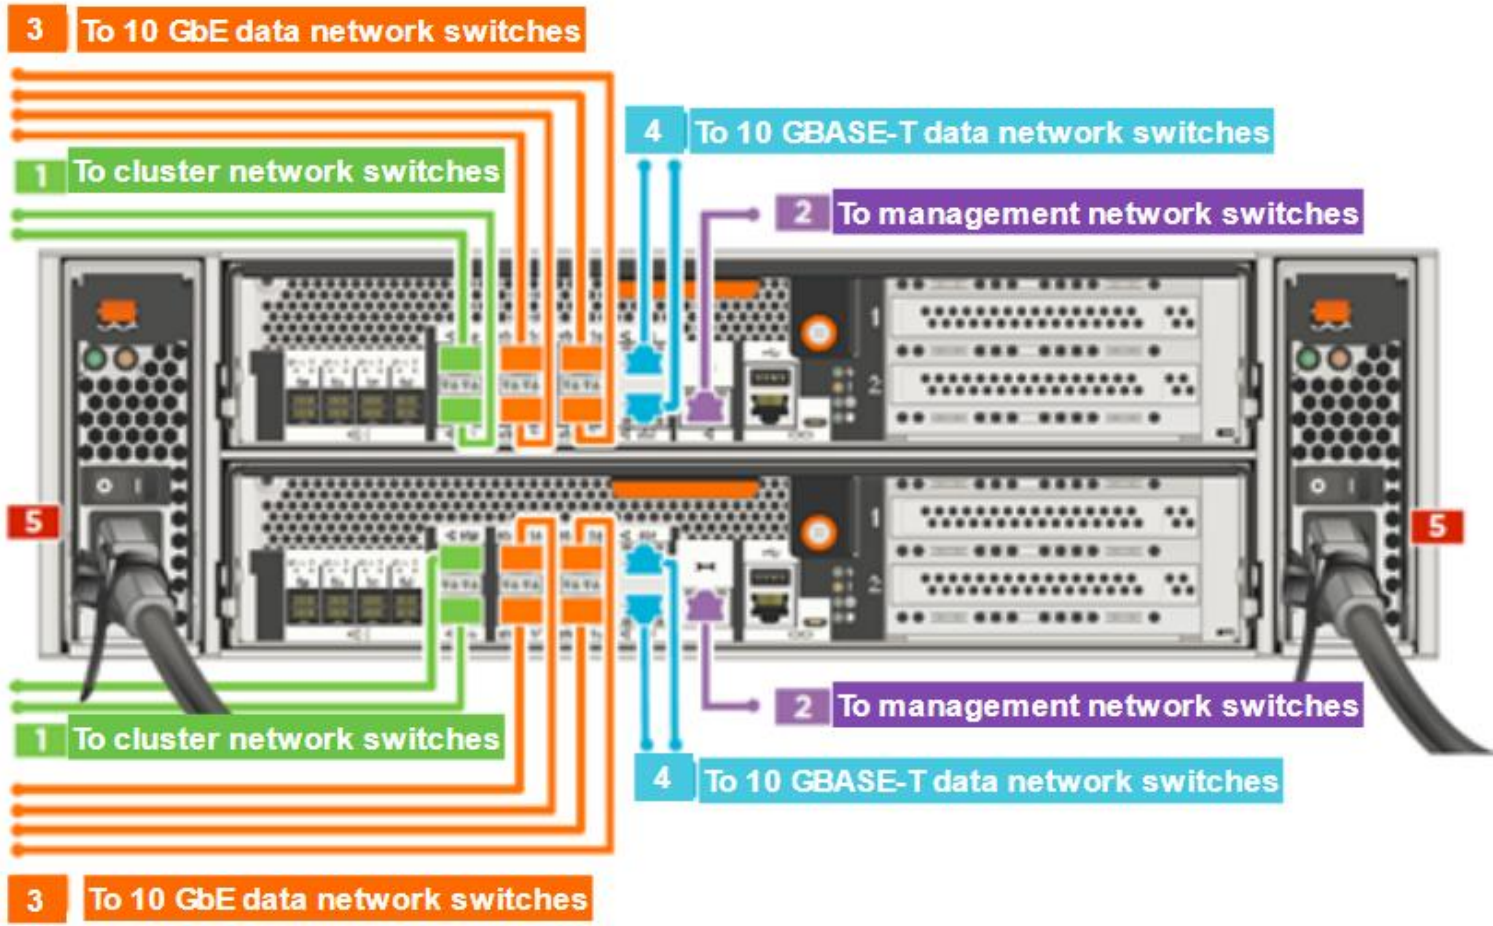

Cabling storage on a two-node switchless cluster with no external storage

Users must cable the SAS ports together in a two-node switchless cluster so that both nodes can access the partner's onboard storage.
Cable the SAS port controller 1 0a to controller 2 0b and controller 1 0b to controller 2 0a.

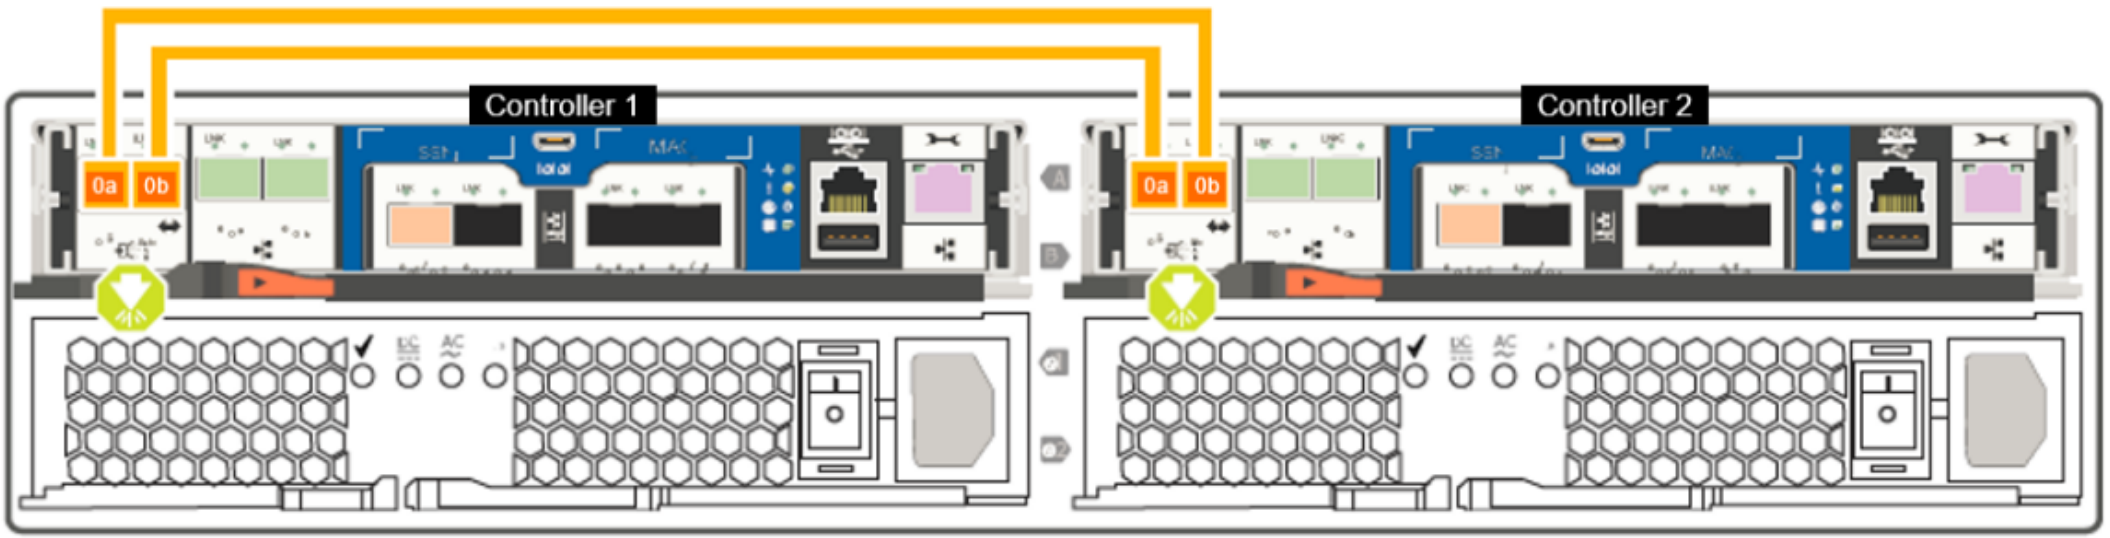

DM3000/5000 series

MultiPath High-Availability cabling rules – DM3000/DM5000 series

- Each expansion enclosure has two I/O modules installed – IOM A and IOM B.
- An enclosure stack is one to 10 enclosures cabled together.
- Cabling the IOM A of an enclosure to the IOM A of another enclosure creates a daisy-chain called the A Path. Connecting B IOMs together creates a B path.
  - The A Path is also referred to as the Primary Path
  - The B Path is also referred to as the Return Path
  - The A and B paths provide communication from the storage controller to the disks and constitute the basis of MultiPath High-Availability (MPHA) cabling.
- A storage controller that 'owns' an enclosure stack connects to the A/Primary Path on the first enclosure and the B/Return Path on the last enclosure in the stack.
- The partner storage controller in an HA pair connects to the B/Return Path on the first enclosure and the A/Primary Path on the last enclosure in the stack.
- Paths load-balance so disks might be accessed over the A or B path.

Single Point Of Failure and MPHA cabling

A Single Point of Failure (SPOF) is a part of a system that, if it fails, will stop the entire system from working. The DM Series storage systems MPHA cablings are designed with redundant hardware components to reduce the single points of failure.

- Disk shelves are connected to each other and to the storage controller with two cables to protect against SPOFs.
- Careful consideration must be given when selecting which ports the cables are connected to on a storage controller to prevent the creation of an SPOF.

MPHA cabling a DM3000/5000 with expansion enclosure

Users must cable the enclosure-to-enclosure connections, and then cable both controllers to the drive shelves. The illustration uses the DM240 as an example. Cabling on the DM120S and DM600S is similar.

1. Cable the enclosure-to-enclosure ports (orange line).
   - Port 3 on IOM A to port 1 on the IOM A on the enclosure directly below
   - Port 3 on IOM B to port 1 on the IOM B on the enclosure directly below
2. Connect each node to IOM A in the stack.
   - Controller 1 port 0b to IOM A port 3 on the last drive enclosure in the stack
   - Controller 2 port 0a to IOM A port 1 on the first drive enclosure in the stack
3. Connect each node to IOM B in the stack.
   - Controller 1 port 0a to IOM B port 1 on the last drive enclosure in the stack
   - Controller 2 port 0b to IOM B port 3 on the first drive enclosure in the stack

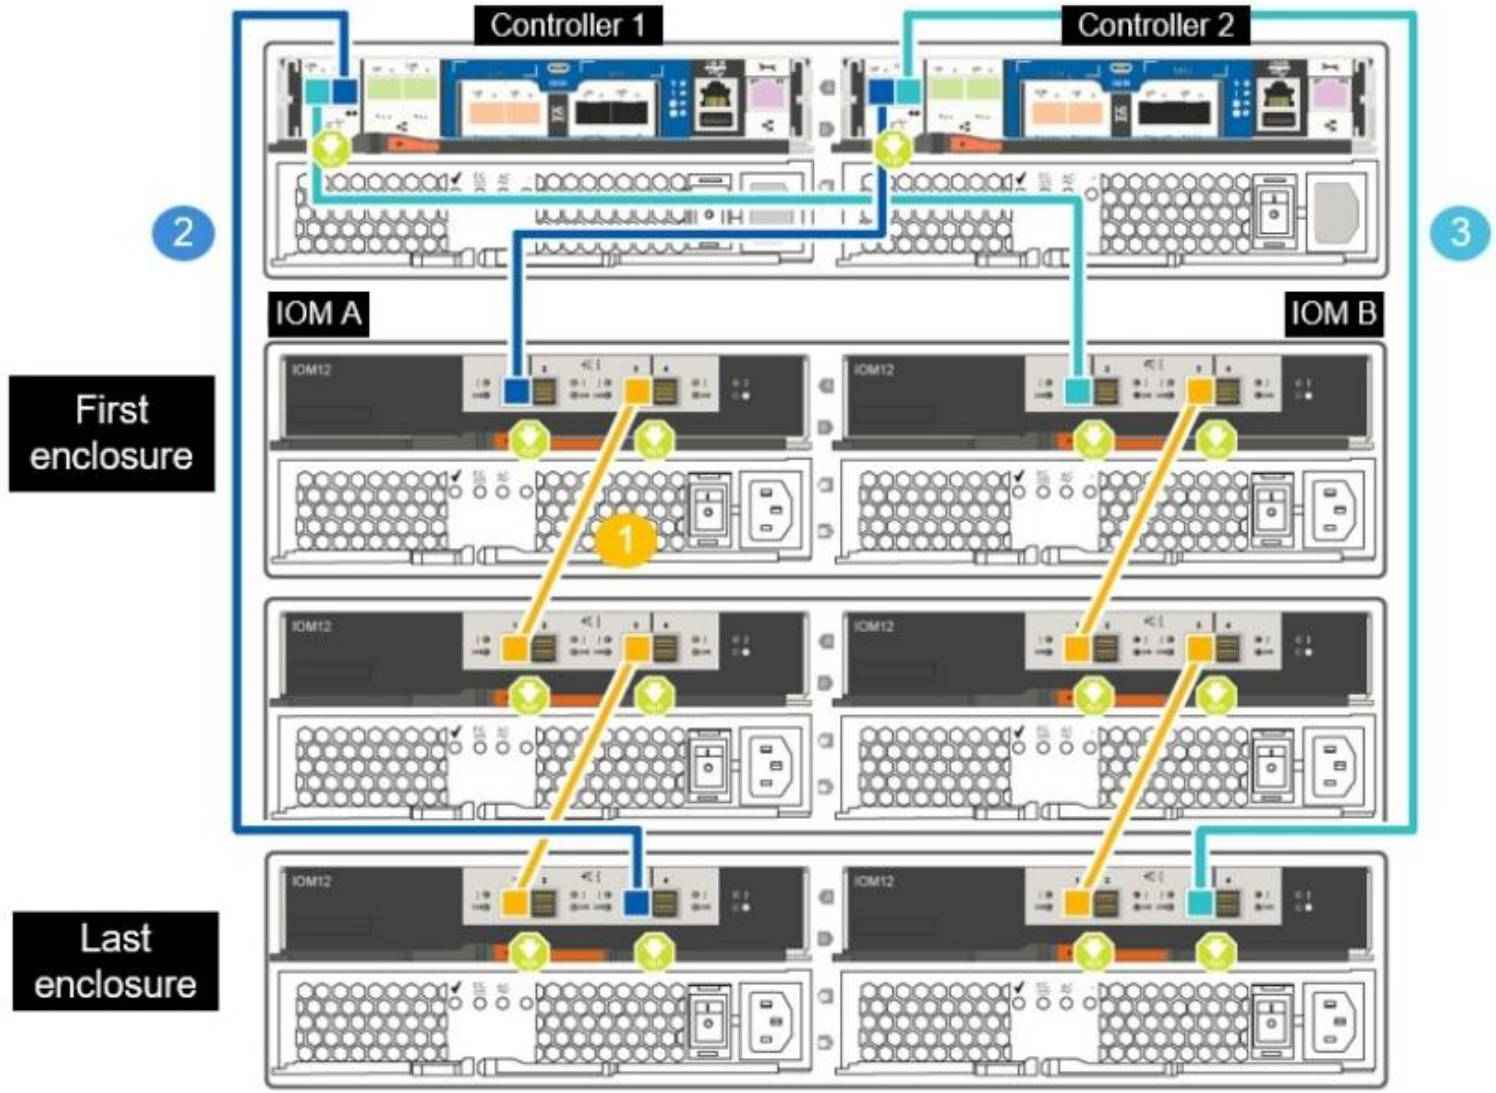

MultiPath High-Availability cabling rules – DM7000 series

Controller A (0a) and C (0c) port cabling rules:

- A and C ports are always the primary paths to a stack.
- A and C ports always connect to the logical first enclosure in a stack.
- Controller one A and C ports always connect to IOM A (domain A).
- Controller two A and C ports always connect to IOM B (domain B).

Controller B (0b) and D (0d) port cabling rules:

- B and D ports are always the secondary paths to a stack.
- B and D ports always connect to the logical last enclosure in a stack.
- Controller one B and D ports always connect to IOM B (domain B).
- Controller two B and D ports always connect to IOM A (domain A).

MPHA cabling a DM7000 series with expansion enclosure

The ThinkSystem DM7000 series supports up to four enclosure stacks with up to 10 enclosures in each stack. For one or two stacks, the integrated SAS expansion ports on the DM7000 controller can be used. For three or four stacks, an additional four-port SAS adapter card is required for the DM7000 controller.
Users must cable the enclosure-to-enclosure connections, and then cable both controllers to the expansion enclosures. The illustration uses the DM240 as an example. Cabling on the DM120S and DM600S is similar.

Click the buttons to see the different cabling illustrations.
The illustration uses the DM240 as an example. Cabling on the DM120S and DM600S is similar.

Two stacks

Four stacks — Controller 1 — Controller 2 — SAS adapter 1 — SAS adapter 2

Lenovo

MPHA cabling a DM7000 with two enclosure stacks

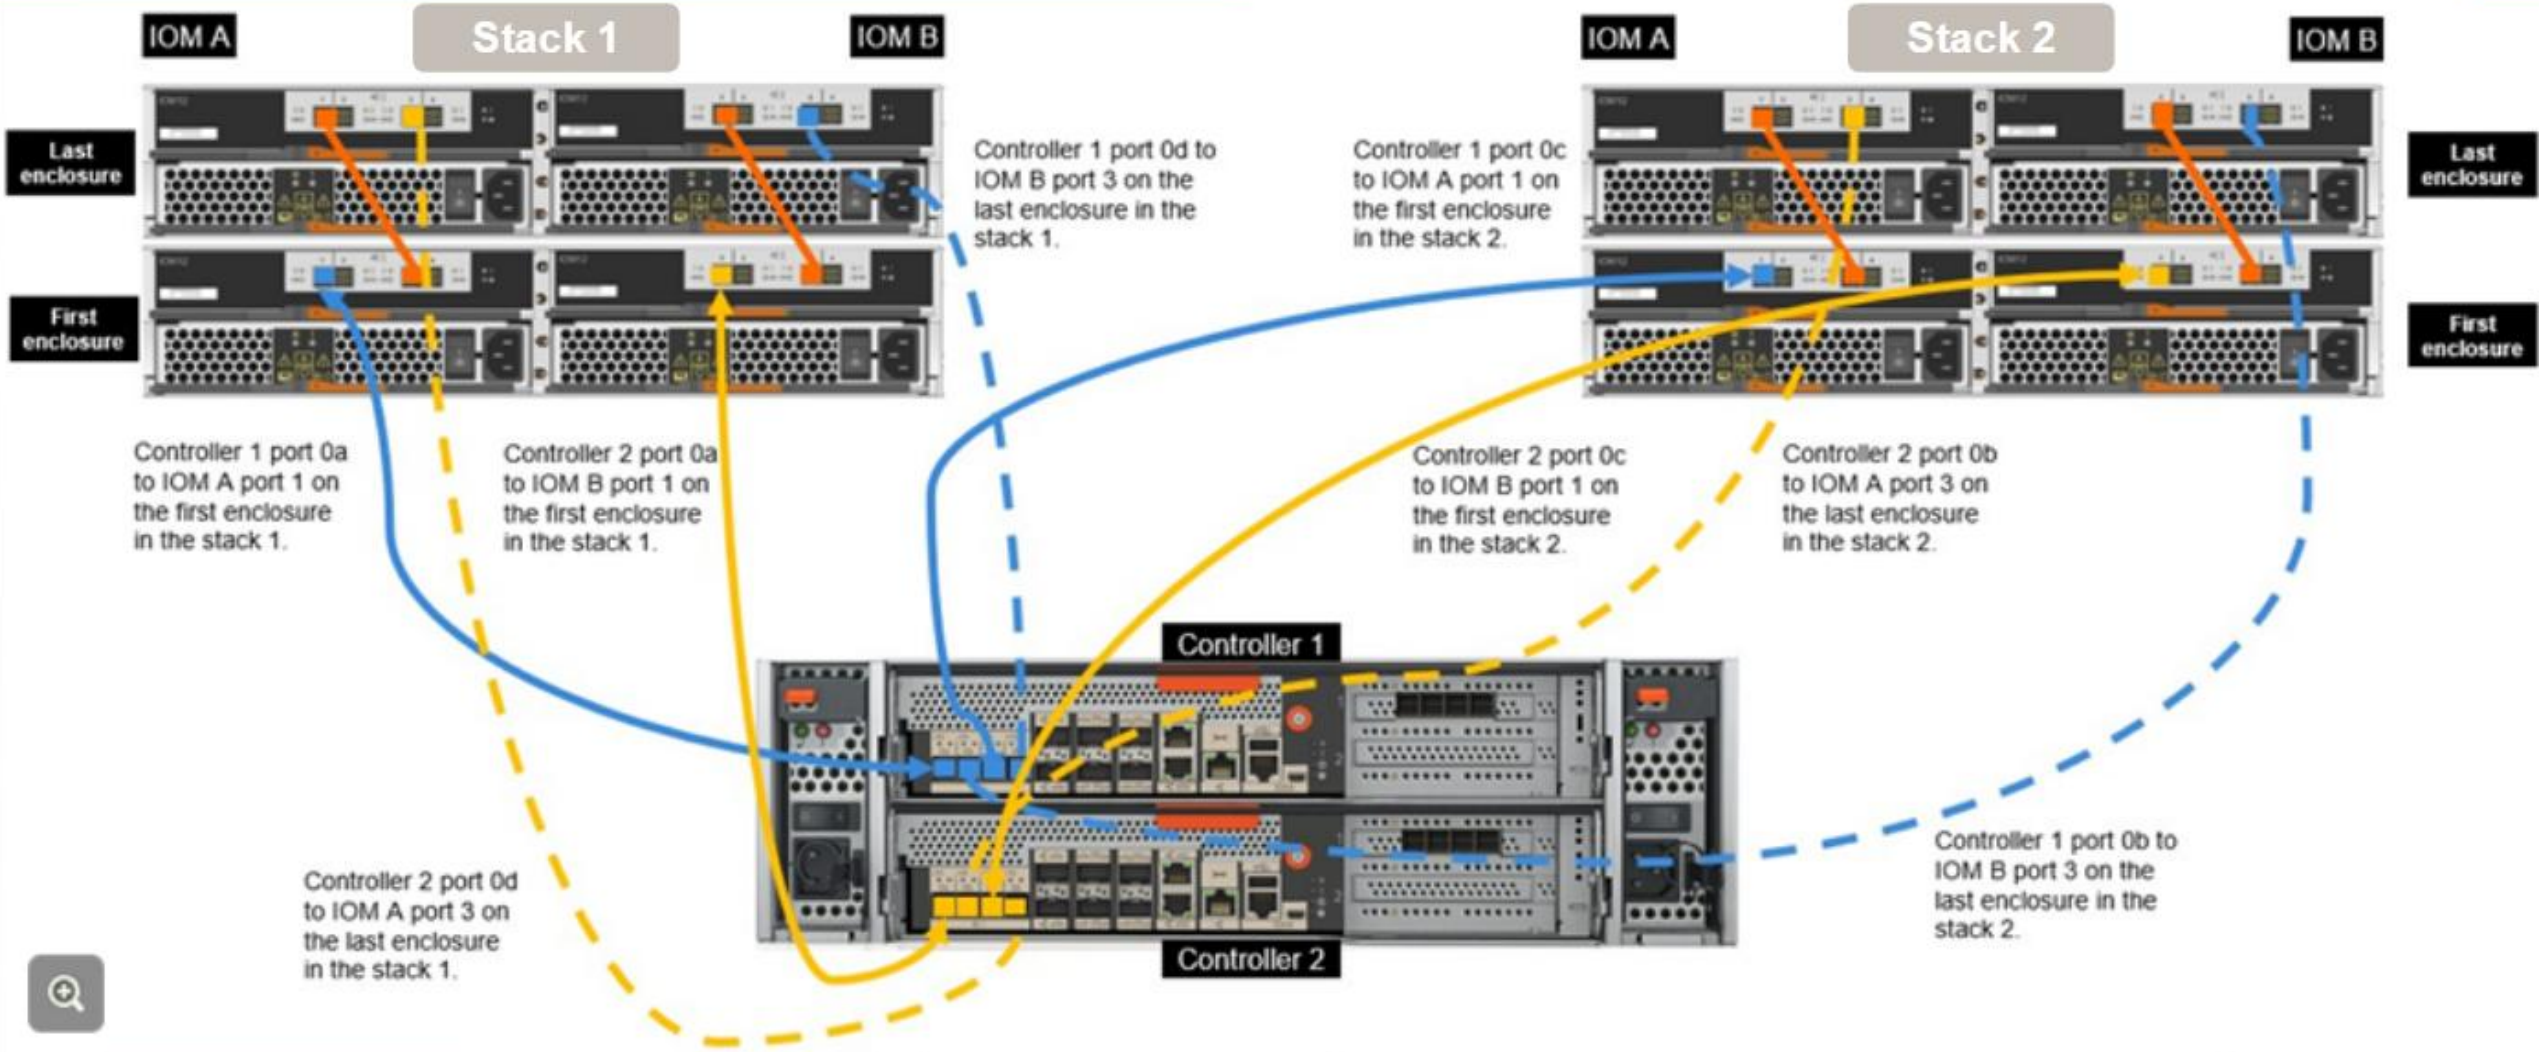

MPHA cabling a DM7000 with four enclosure stacks

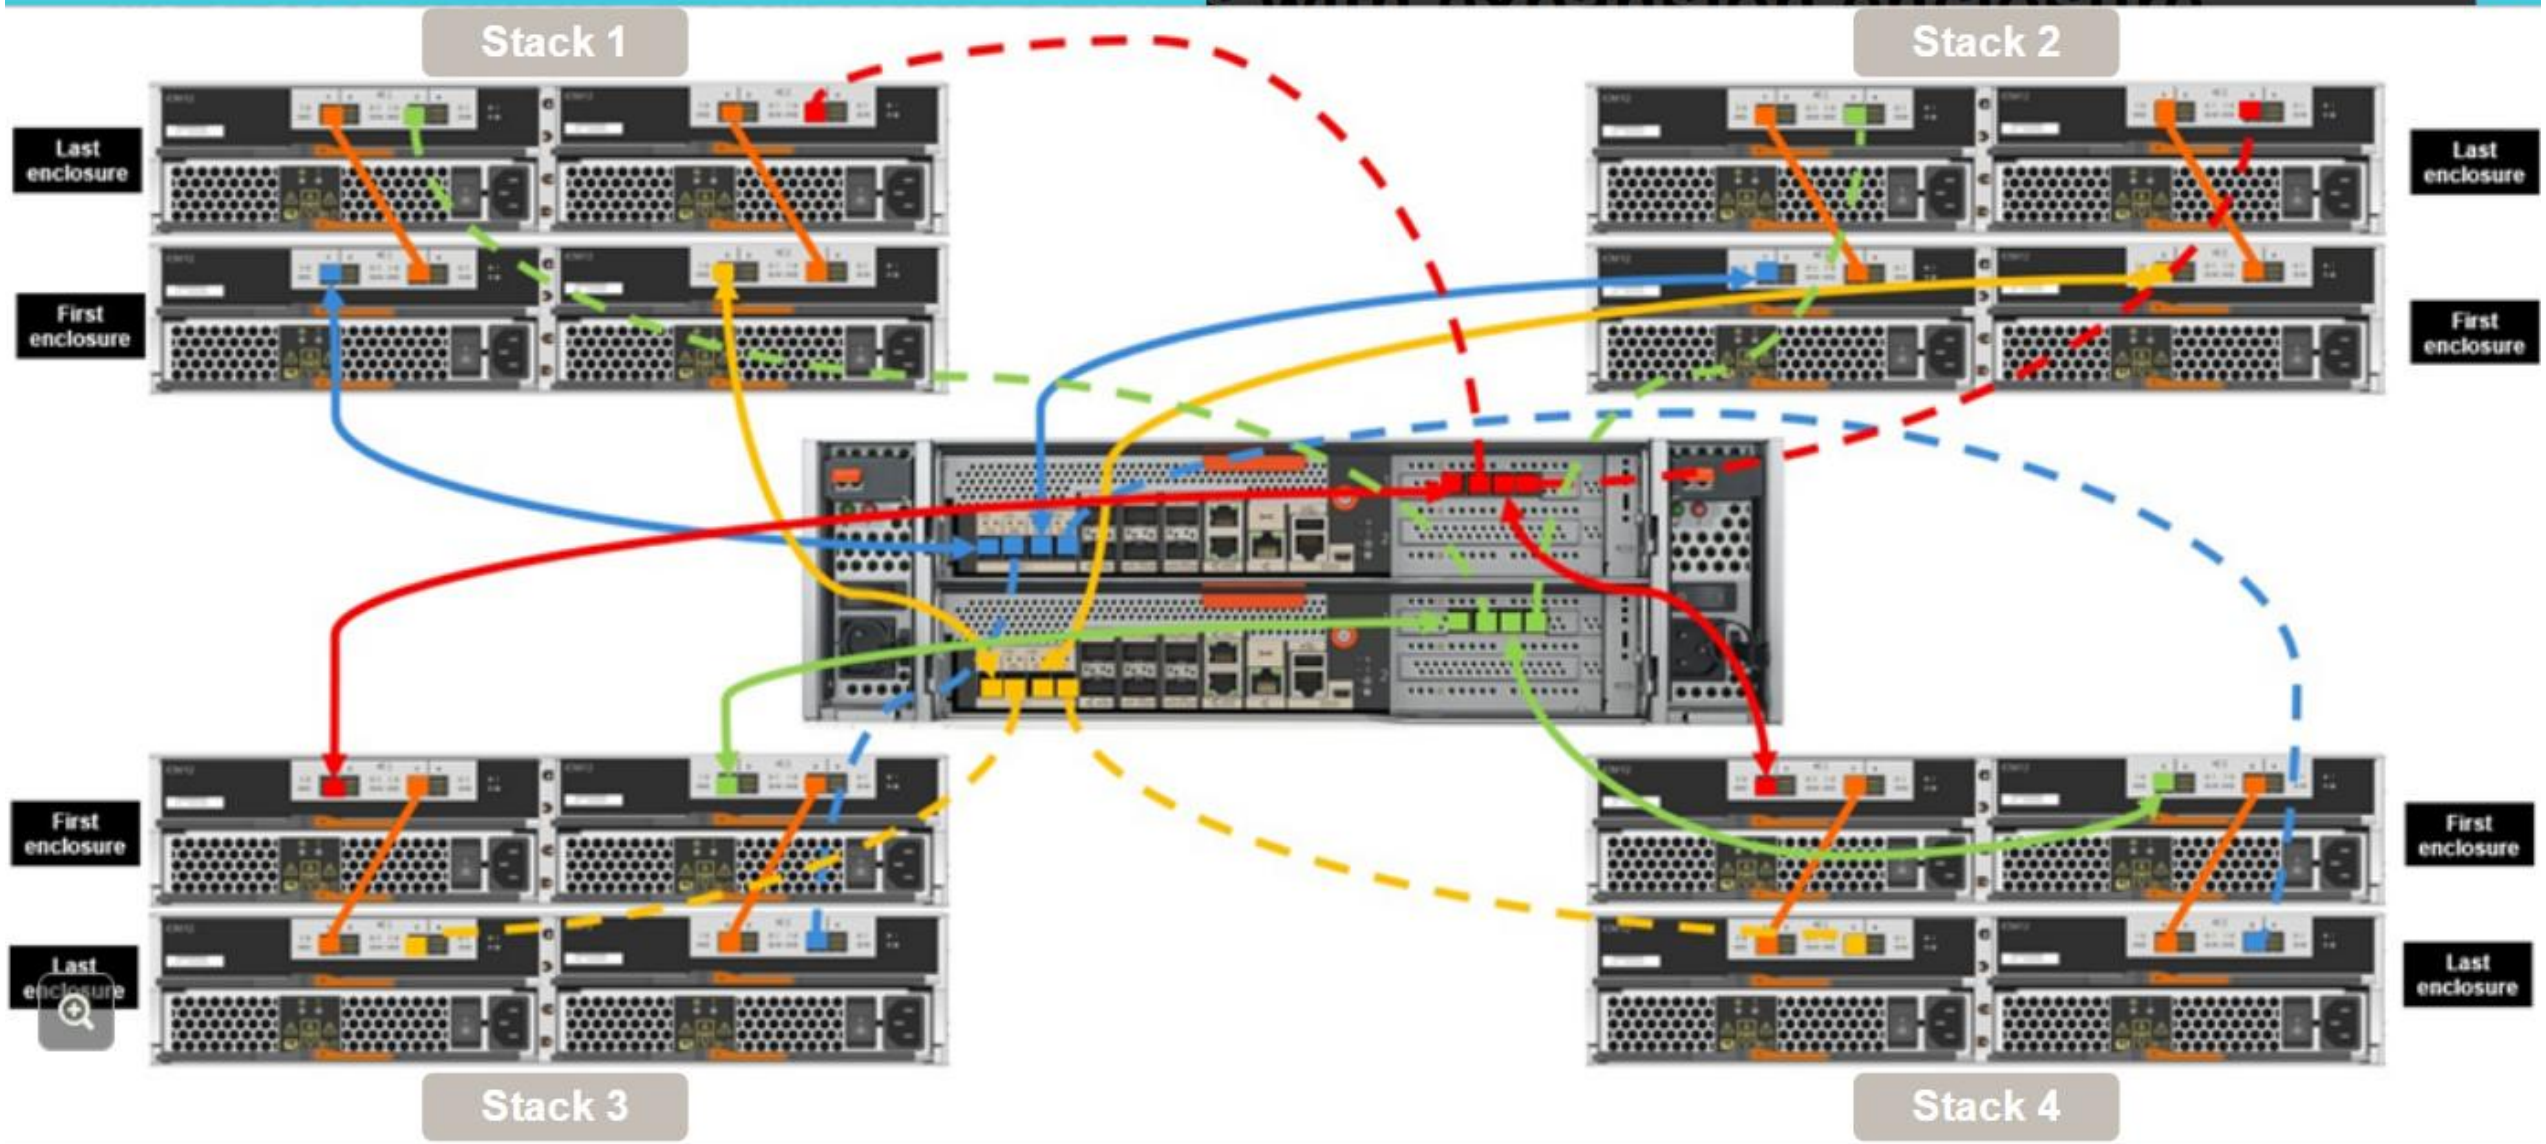

Controller 1 cabling diagram in a four-stack configuration

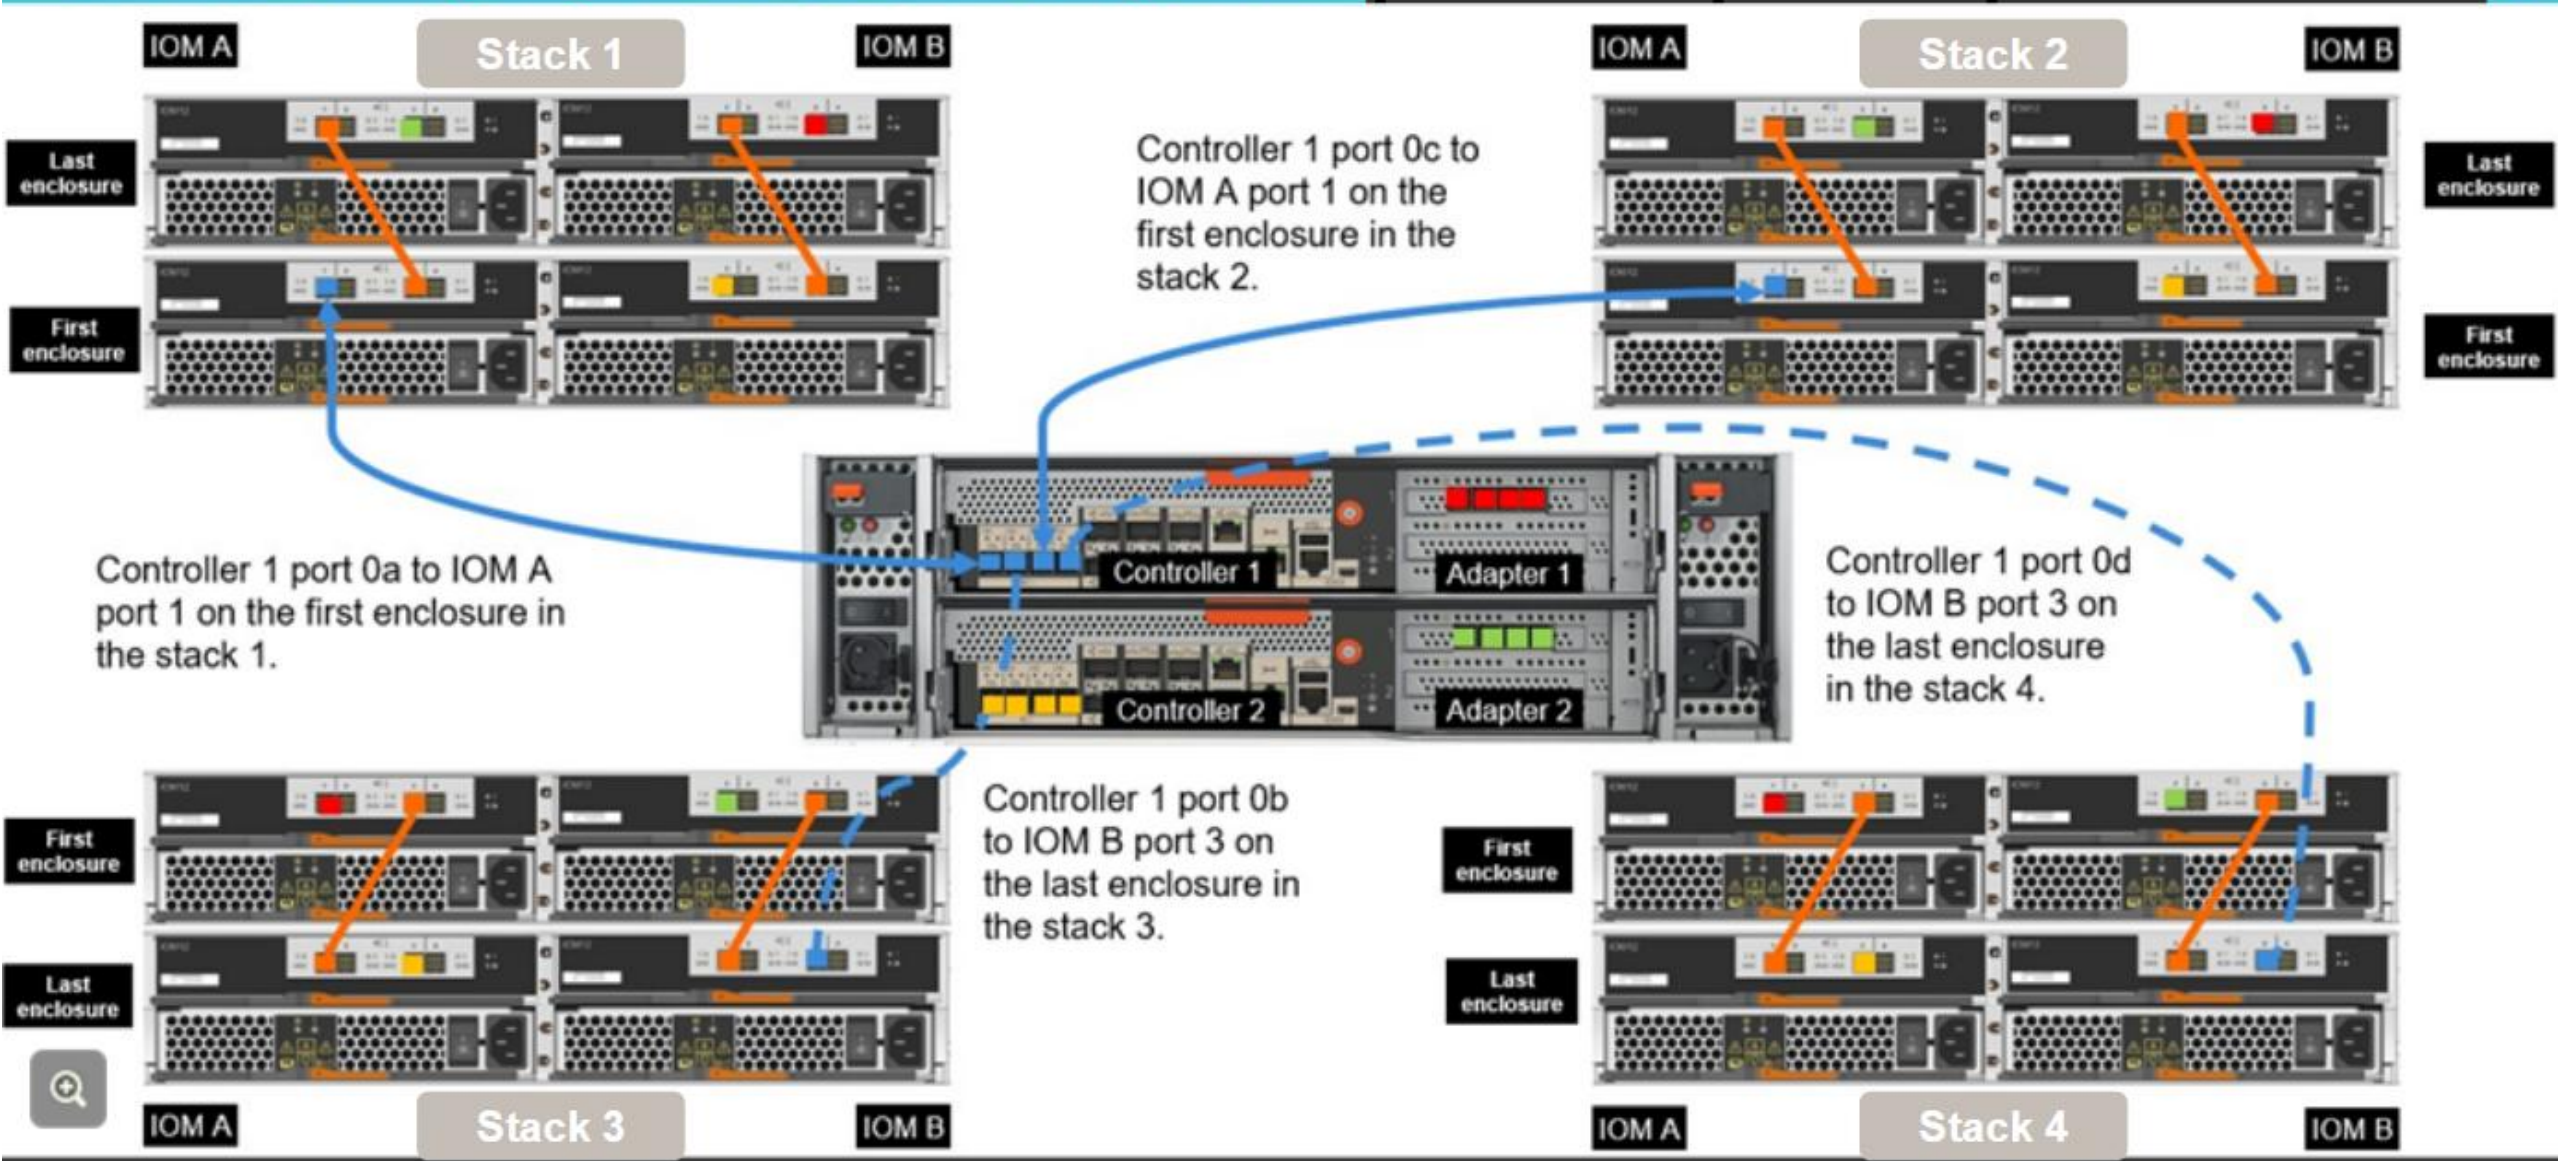

Controller 2 cabling diagram in a four-stack configuration

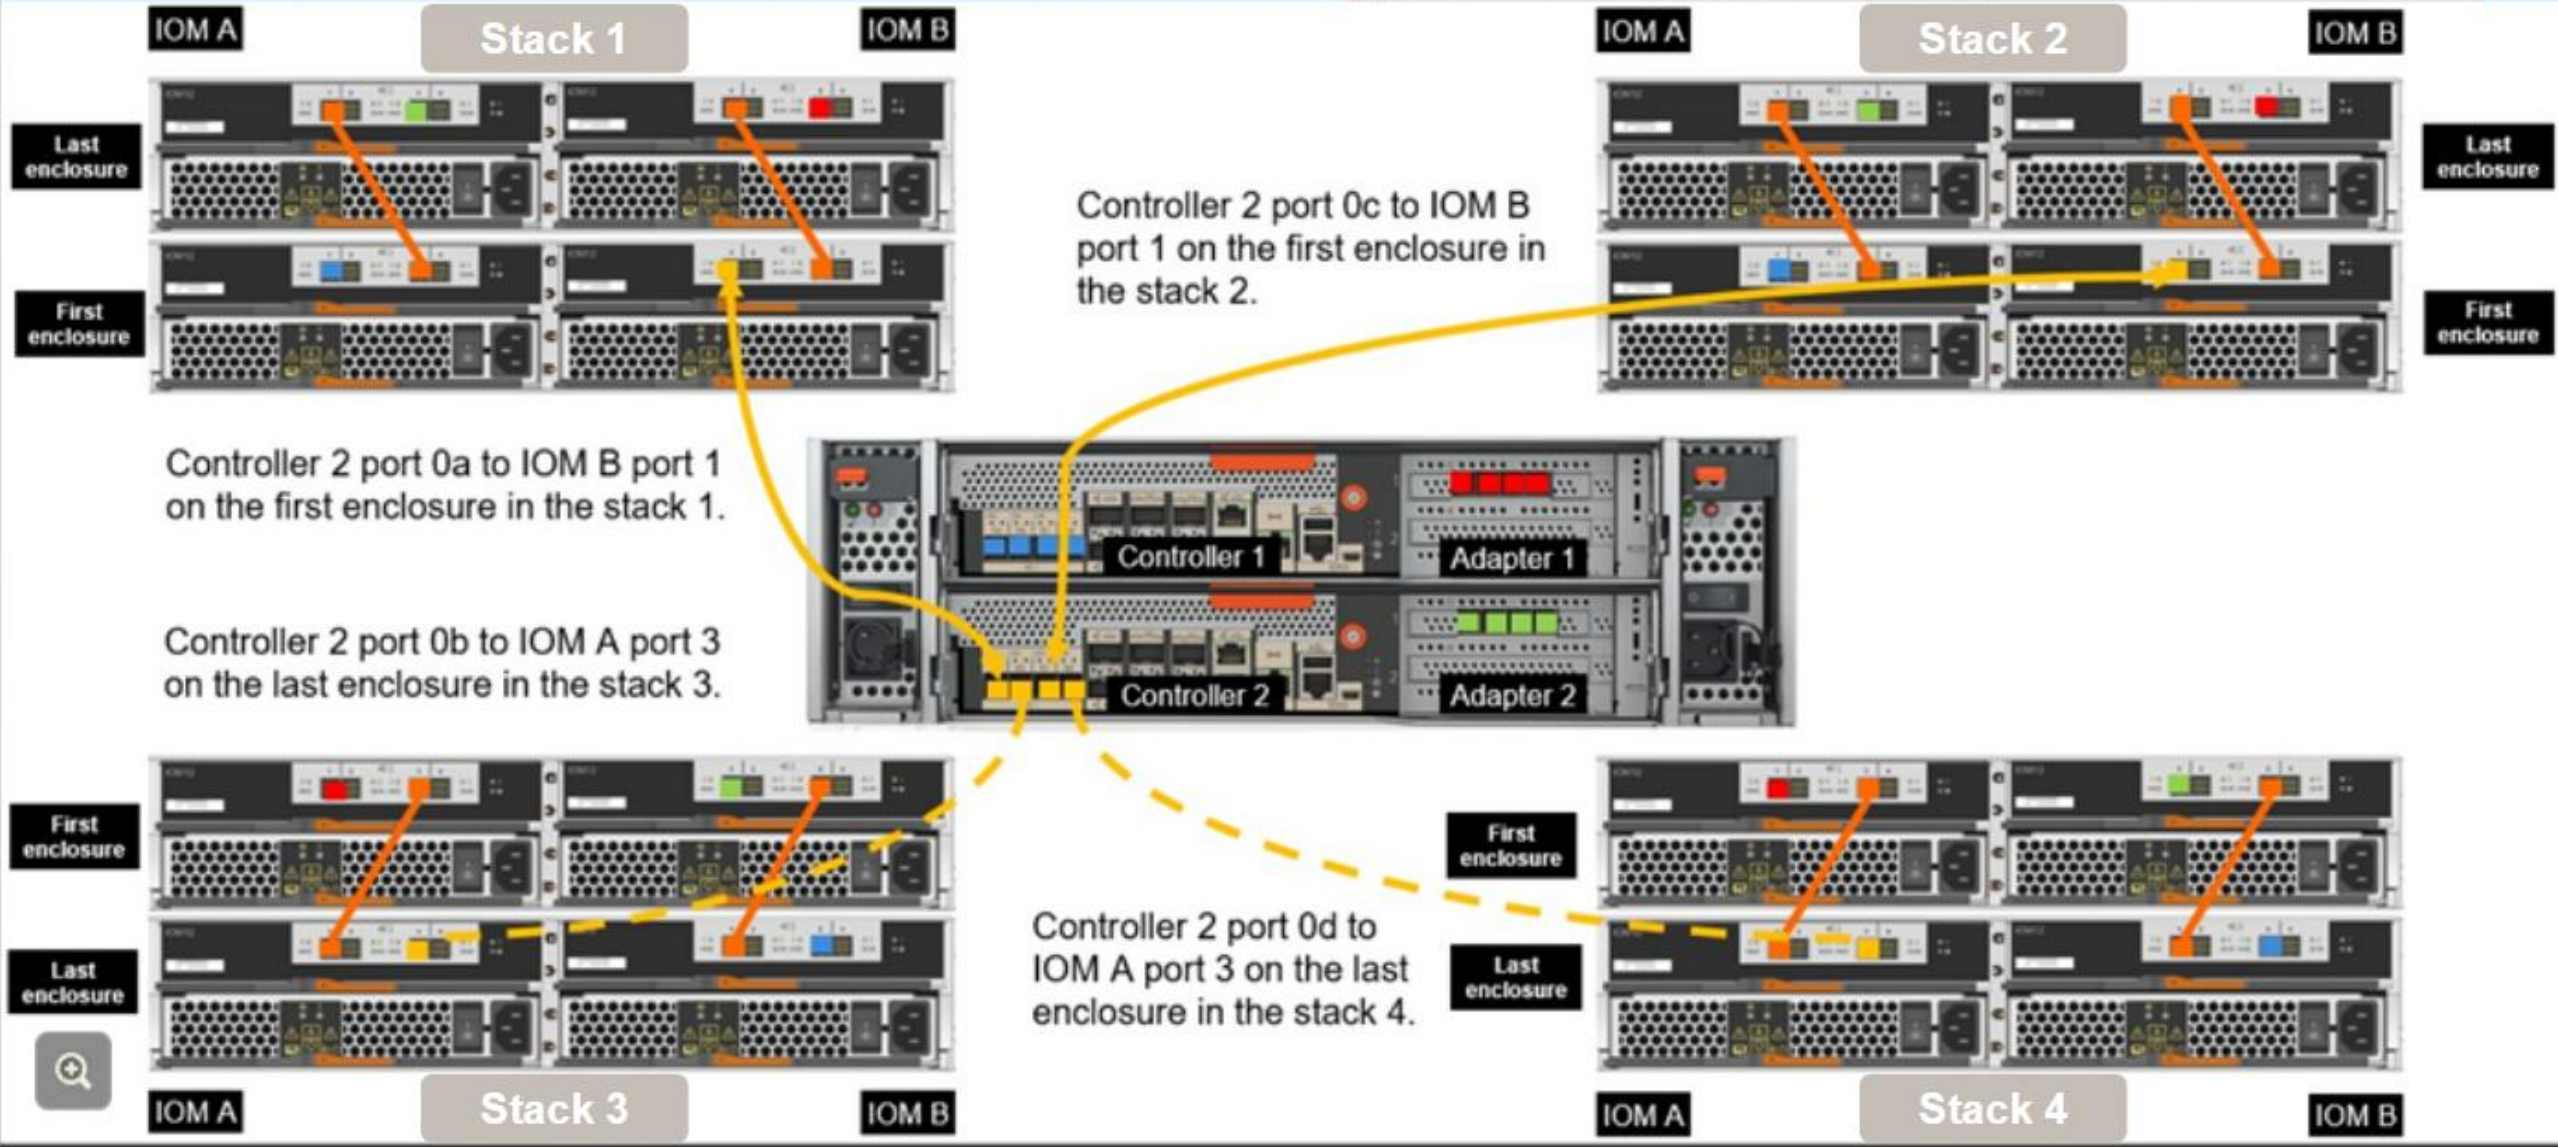

SAS adapter 1 cabling diagram in a four-stack configuration

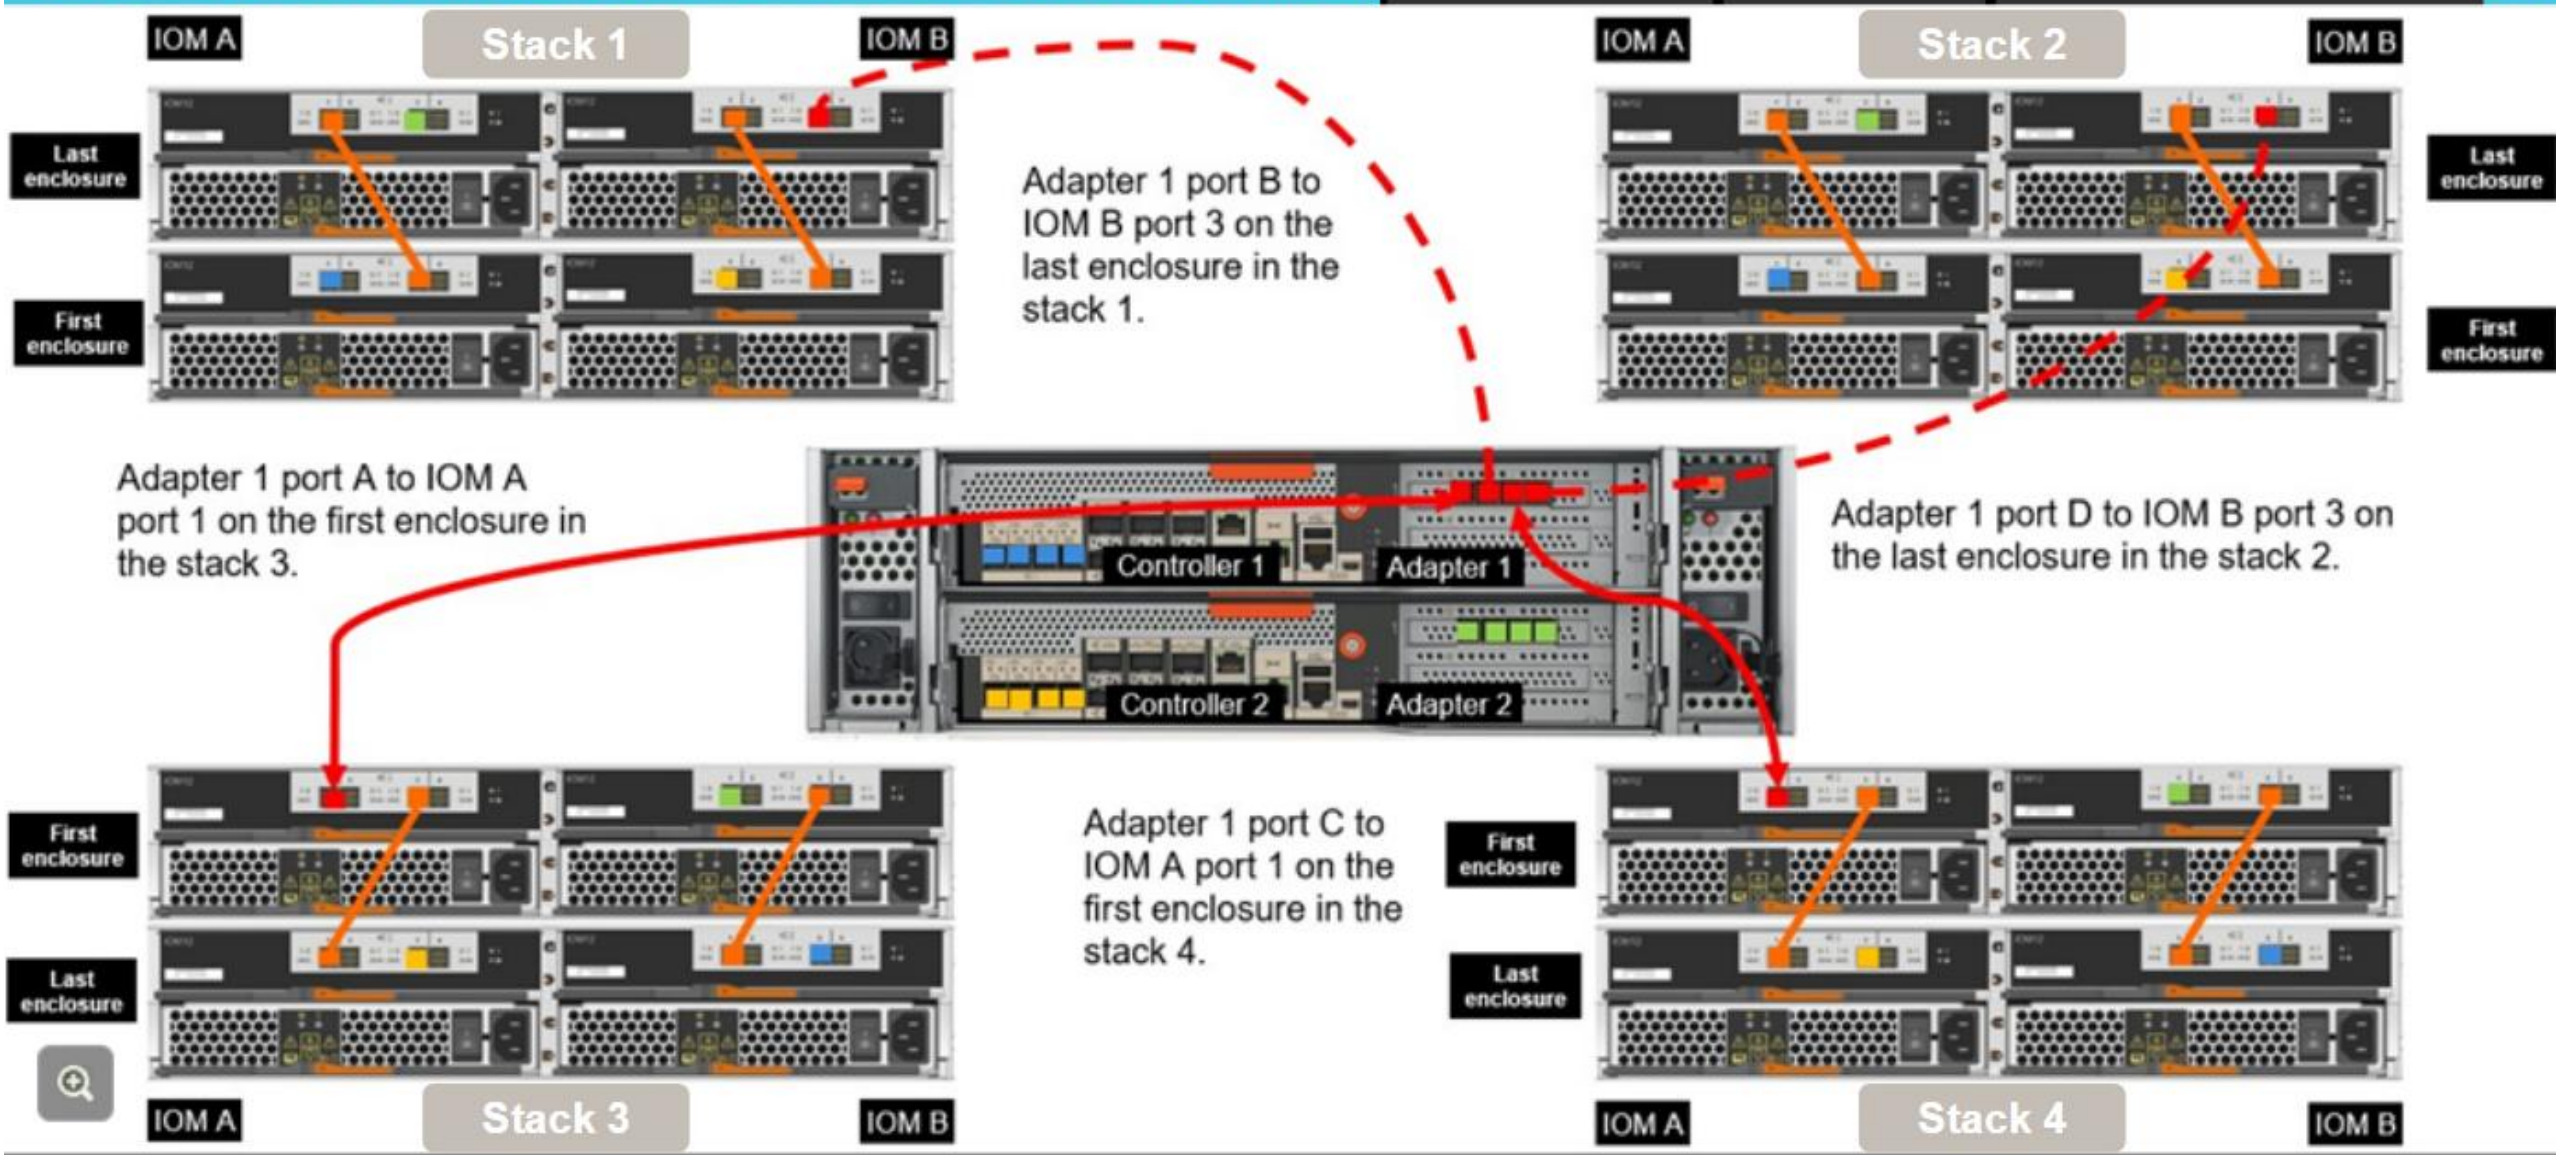

SAS adapter 2 cabling diagram in a four-stack configuration

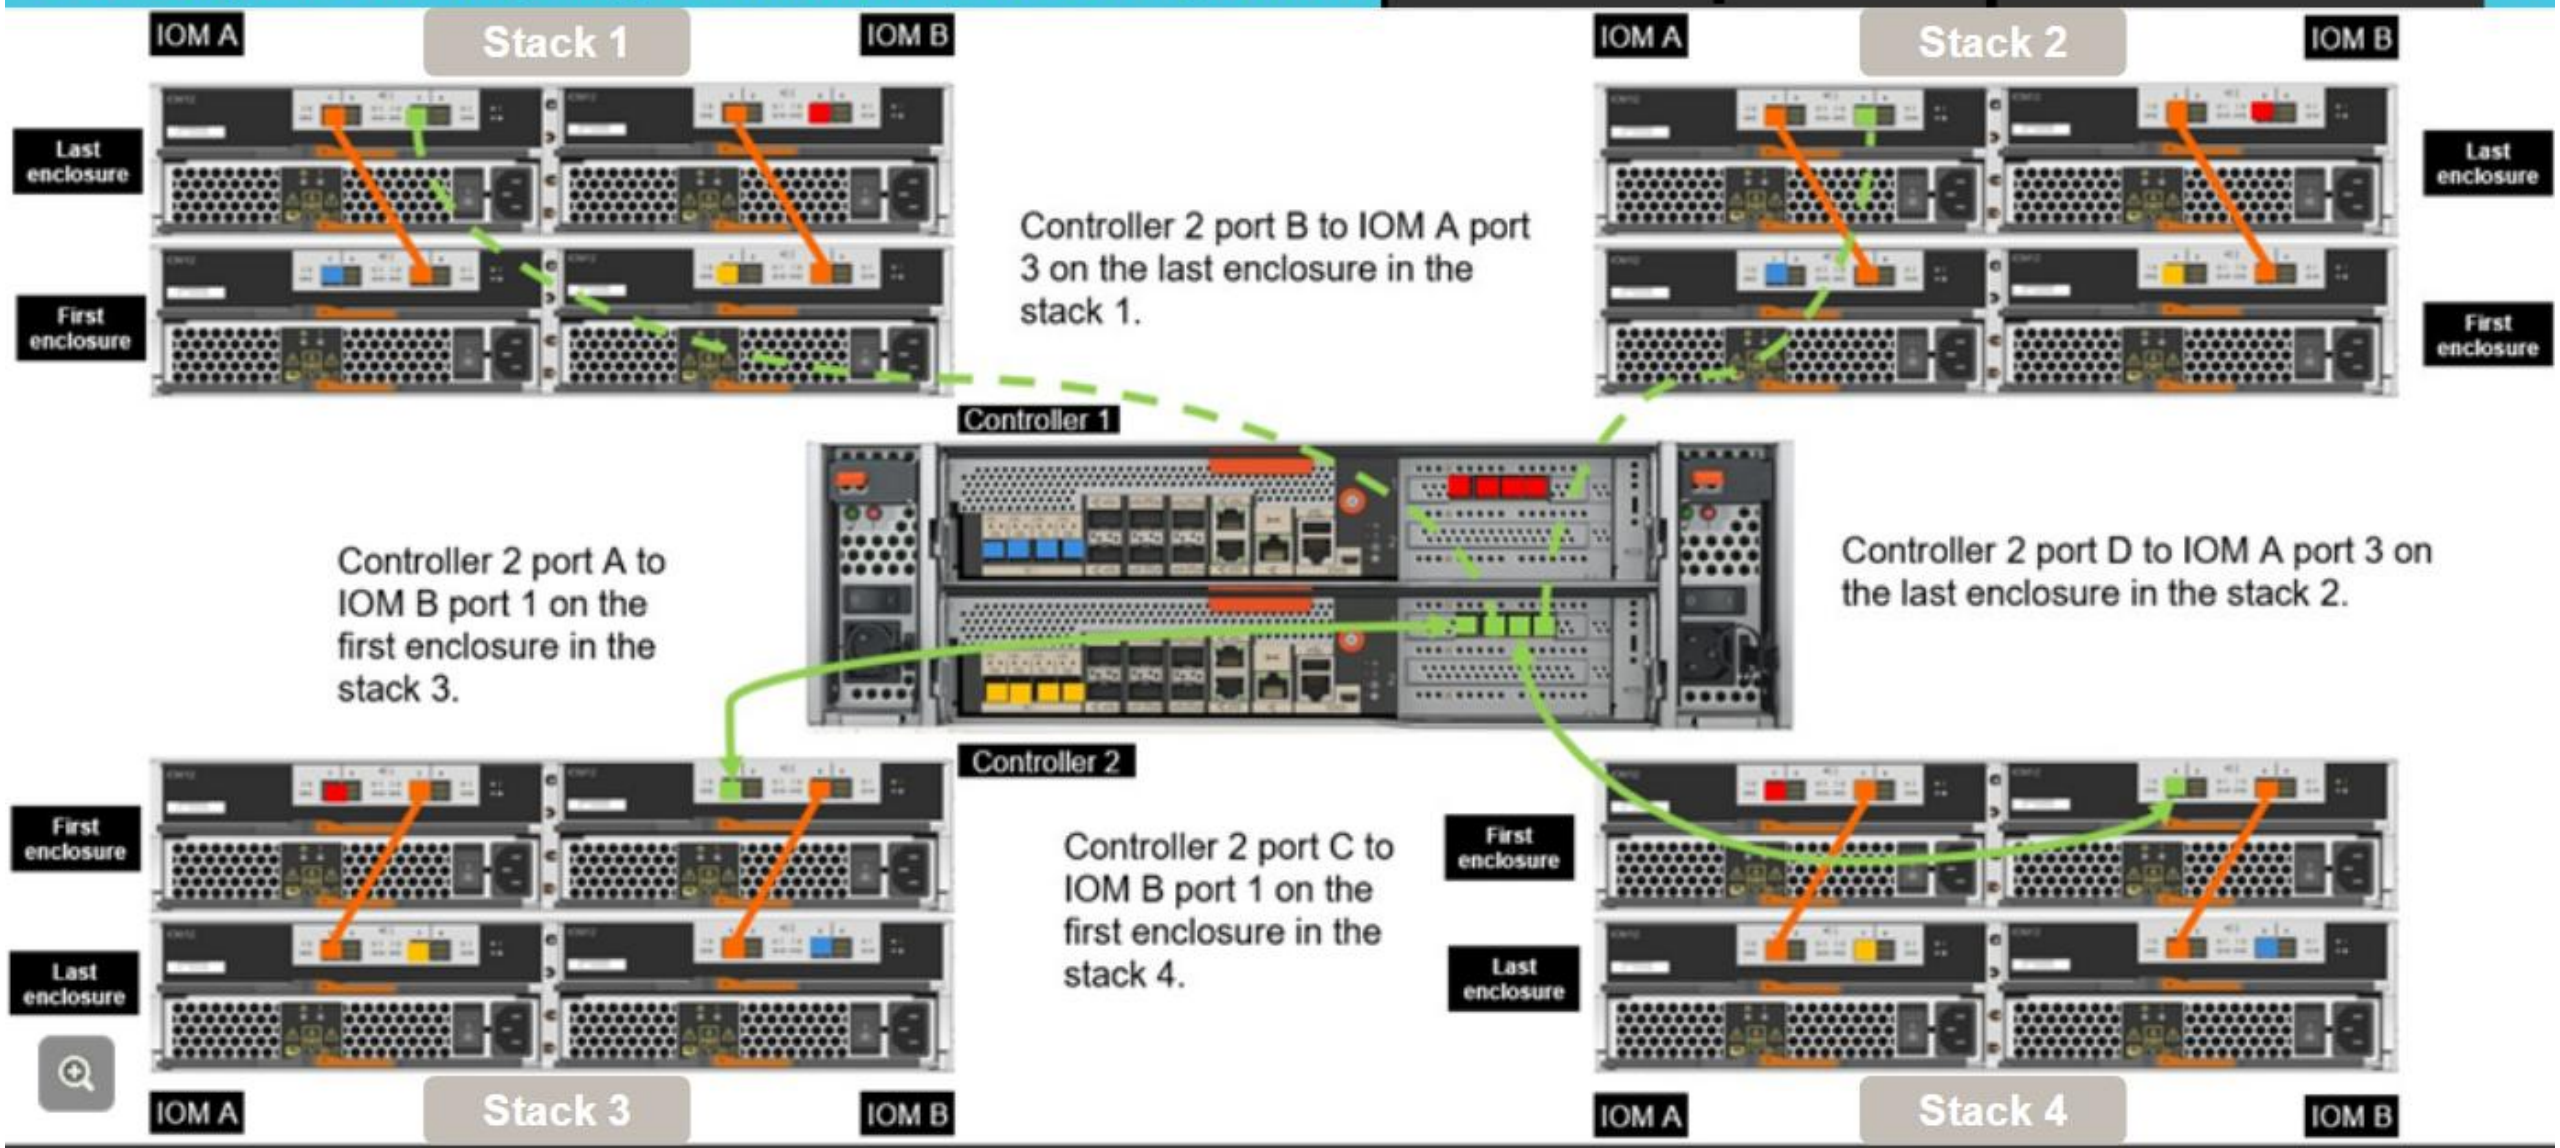

Cabling and configuring your client

After cabling the controller and expansion enclosure, connect to the console port and management port of the controller to complete the system setup and configuration.

- Connect the console cable to the client and the console port on the controller module (mini-USB or RJ-45 console port).
- Connect the client to the switch on the management subnet.
- Assign a TCP/IP address to the client – use one that is on the management subnet.
- Set the console port on the client to the following setting:
  - Baud rate: 115200
  - Data bits: 8
  - Parity: None
  - Stop bits: 1

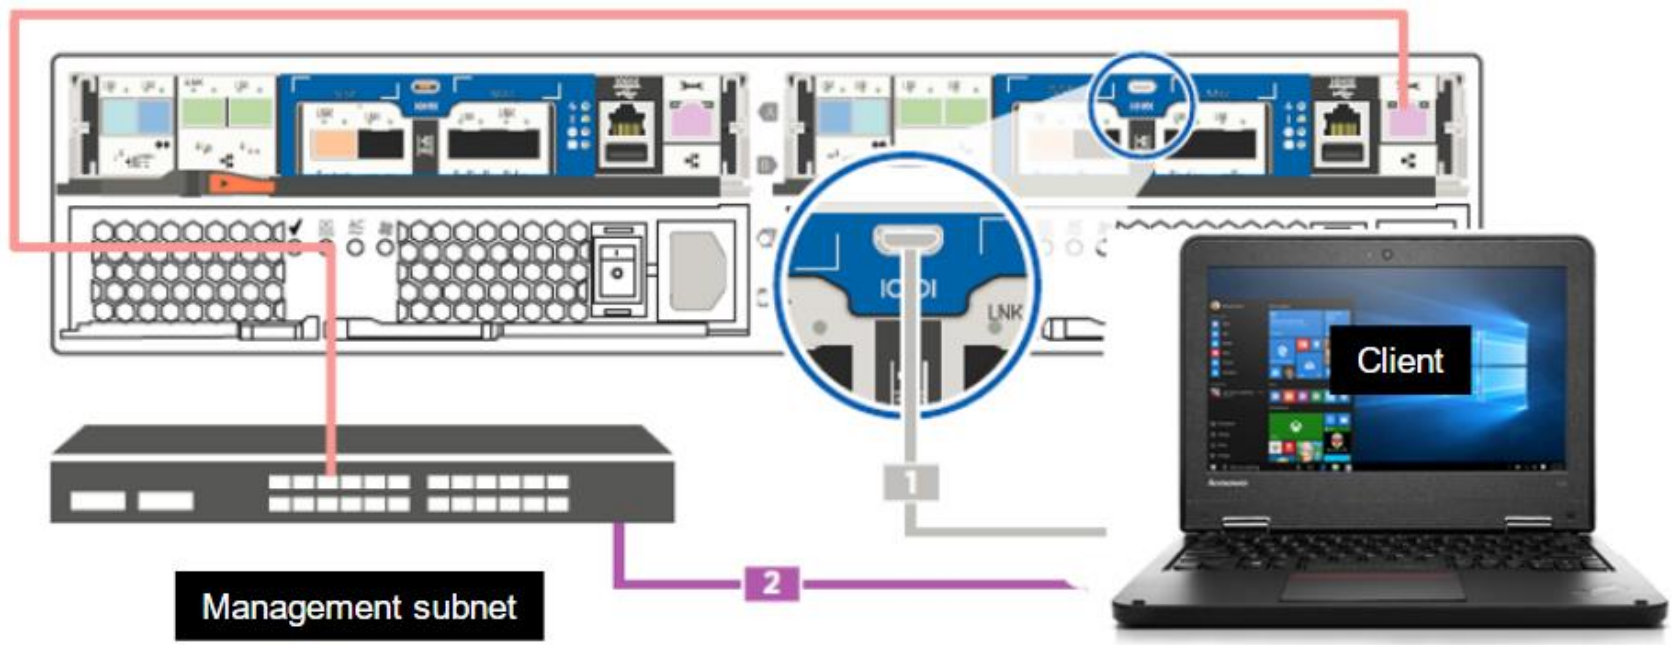

Setting the expansion enclosure ID

If the system has one or more drive shelves, set the enclosure IDs:

1. Power on the drive enclosure, and then remove the front bezel.
2. Press and hold the red button until the first digit flashes, and then press the button to advance the first digit (0-9) to the desired number. The first digit continues to flash.
3. Press and hold the button until the second digit flashes, and then press the button to advance the second digit (0-9) to the desired number. The first digit stops flashing, and the second digit continues to flash.
4. Press and hold the button until the second digit stops flashing, and then replace the end cap on the enclosure.
5. Wait about 10 seconds for both digits to start flashing again and for the LED to light, and then power-cycle the drive enclosure to make the enclosure ID take effect.
6. Repeat these steps for any remaining drive shelves.

For more details, refer to the following video link on how to configure the SAS enclosure ID – start from 2:49.
https://www.youtube.com/watch?v=yn2Z4b0ACTs

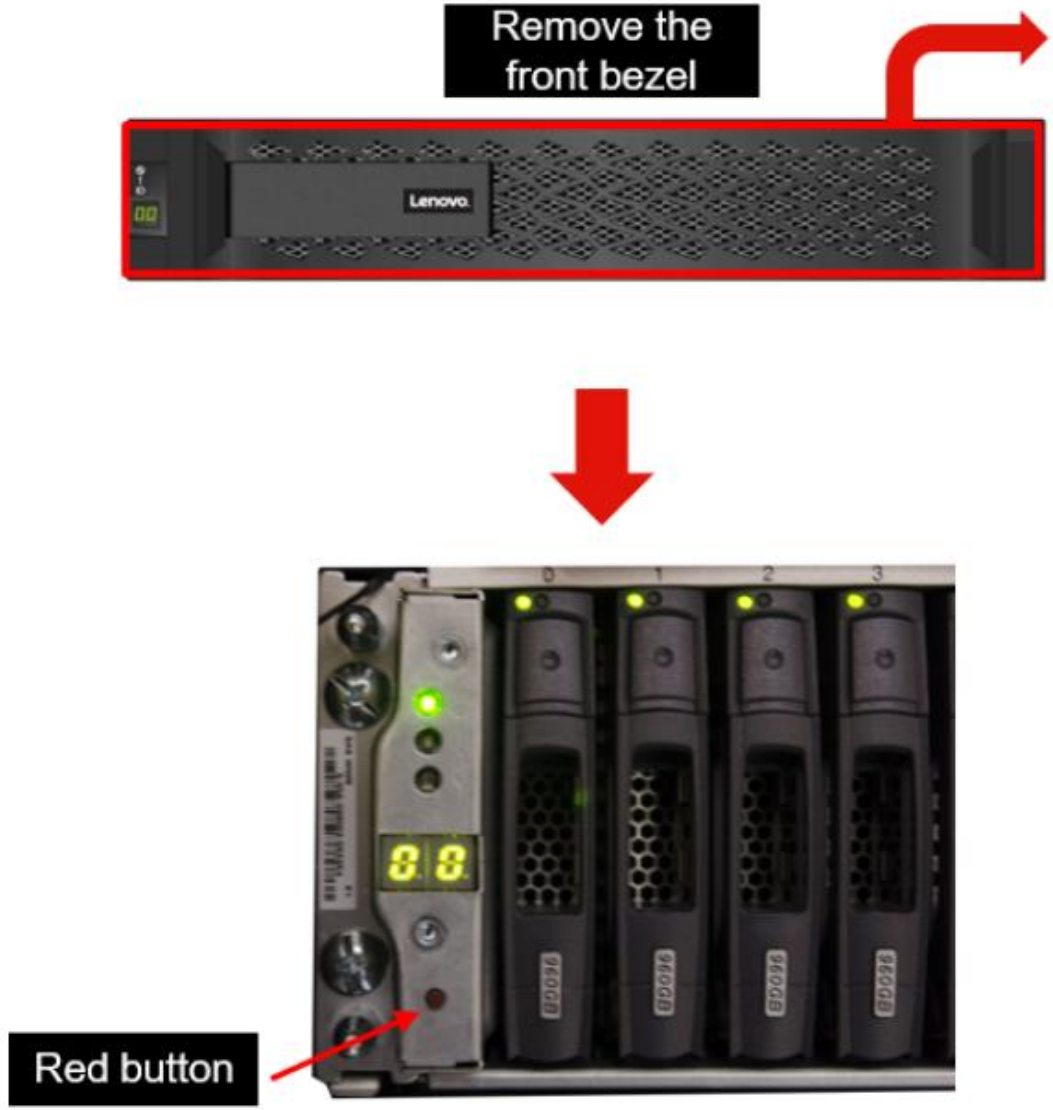

Powering on the system for the first time and completing the settings

To power on the system (assuming that all the connected expansion enclosure IDs have been set and are already powered on):

- Plug the power cords into the power supplies of the controller, and then connect them to power sources on different circuits.
- Turn on the power switches on both nodes.

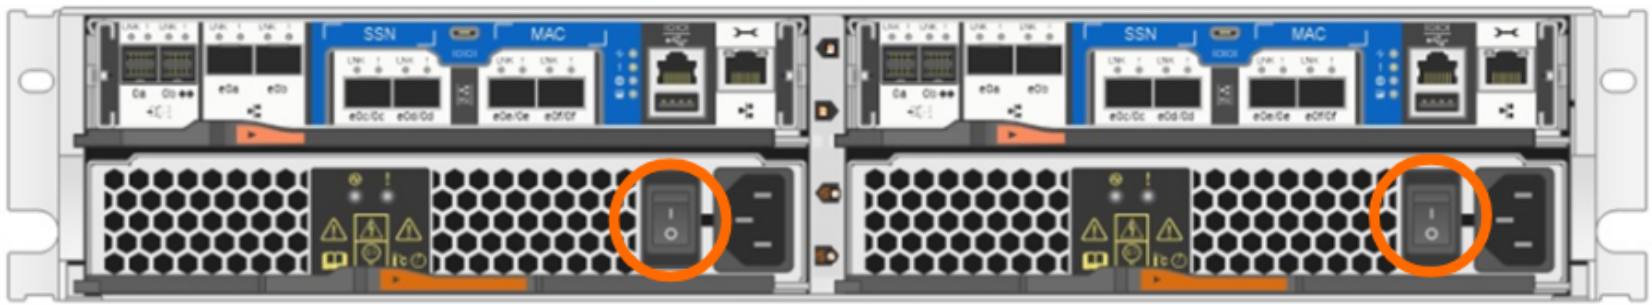

Assign an initial node management IP address to one of the nodes. If the user's environment has DHCP configured on the management network, record the IP address assigned to the new controllers. Otherwise:

1. Open a console session using PuTTY, a terminal server, or the equivalent for your environment.
2. Enter the node management IP address when prompted.

Then, open the browser to the node management IP address, and use the ThinkSystem Storage Manager Web GUI to complete the system settings. The format for the address is https://x.x.x.x.

Power on/off procedure

- Powering on the storage system:

1. Turn on the power switches of all the connected DM Series disk shelves.
2. Wait approximately two minutes.
3. Turn on the power switches of DM Series controllers.
4. Issue the `cluster show` command to check the system health.
5. Issue the following commands if the cluster HA or storage failover were disabled:
   - `::>cluster ha modify -configured true`
   - `::>storage failover modify -node * -enabled true`

- Powering off the storage system:

1. Disconnect all connected CIFDS/NFS clients.
2. Offline hosts use FCP or iSCSI-based LUNs before shutting down the storage system.
3. Log in to one node, and then issue the following command:
   - `::>halt !local -inhibit-takeover true`
   - Each non-local node takes several minutes to shut down.
4. Issue the following command on the local node:
   - `::> halt local -inhibit-takeover true`
5. Turn off the power switches of the DM Series controllers.
6. Wait approximately two minutes.
7. Turn off the power switches of all the connected DM Series disk shelves.

```
a300::> cluster show
Node                  Health  Eligibility
--------------------- ------- ------------
a300-1                true    true
a300-2                true    true
2 entries were displayed.
```

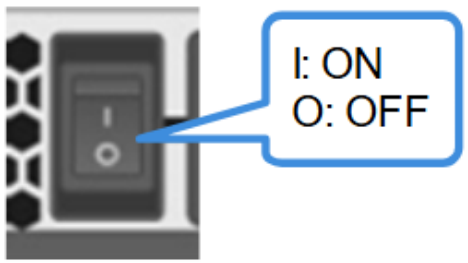

Note: There is no way to power off the system from the systems management interface or the CLI. Users must use the power switches next to the PSU.