GUIDE

HPC Reference Architecture

UTILIZING ROCKPORT'S SCALABLE AUTONOMOUS NETWORK

Contents

HPC Reference Architecture

Introduction

This reference design document helps network designers and system administrators visualize and plan the deployment of the Rockport scale fabric in their data center. It includes a brief review of the necessary High Performance Computing (HPC) components, the approach to designing the Rockport scale fabric, a detailed review of example scenarios, and a network deployment overview.

NOTE

Before reading this document, you should be familiar with the key concepts and underlying technology of the distributed architecture. See the *Rockport Scale Fabric Technology Primer* for reference. If you'd like to know more about the installation process or day-to-day operations, see the *Installation Guide* and the *Administration Guide*, respectively.

Solution overview

At the core of the Rockport scale fabric is the Rockport Network Operating System (rNOS) software, which enables advanced switching without traditional switches. Each server has a Rockport NC1225 Fabric Card, which connects to a passive optical interconnect device (a Rockport SHFL) using a standard MTP/MPO-24 fiber optic cable, allowing for the direct interconnection of the fabric cards (referred to as nodes). Network administrators can monitor the network using the Rockport Autonomous Network Manager (ANM) software.

Rockport's distributed architecture integrates into existing data centers and traditional networks. As illustrated in Figure 1, Rockport network cards are installed in the servers and storage units, interconnecting them through the Rockport SHFL, shown at the top of the rack. This means that the key elements of the HPC system - the compute, storage, cluster administration, login, and management nodes - are all directly interconnected within the Rockport scale fabric.

Administration of the scale fabric is managed by the ANM, which provides visibility into the network, enables troubleshooting, and allows for in-field updating.

In a traditional switch-based architecture, the radix of the switch can limit the size of the compute unit (total number of compute nodes) and often influences the number of nodes per rack. With the Rockport scale fabric, the actual needs of the workloads drive the total number of compute nodes, and the capacity of the rack (power, cooling, space) determines the number of nodes per rack.

Figure 1. Solution overview

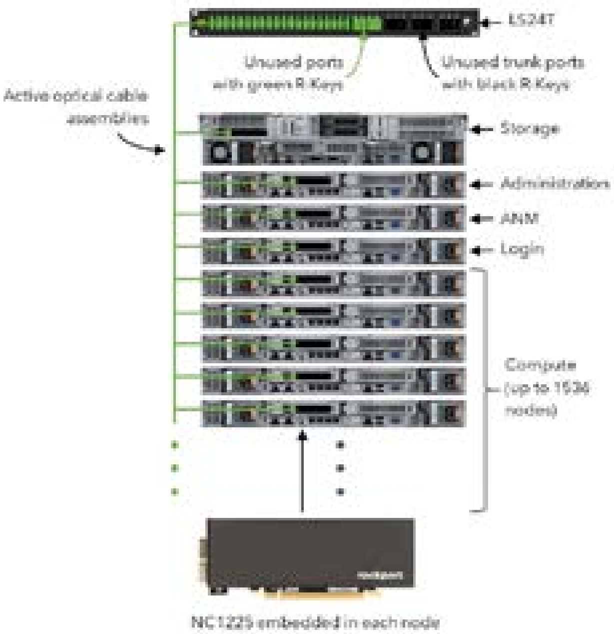

NOTE

Actual power and cooling calculations are out of the scope of this document. Contact your Cerio representative for more information on power and cooling.

Most HPC systems have at least two network fabrics. One is used for administration and management and the second for compute node communication and storage traffic. This document focuses on the network fabric for the compute and storage traffic. Depending on the circumstances, storage may be on its own third network, such as an existing storage network shared by multiple compute clusters. In this case, a gateway can be used to access the network (described later in this document). One benefit of a Rockport scale fabric, however, is that its performance characteristics eliminate the need for a separate storage network, allowing for a converged high-performance fabric with lower operating costs.

Designing a Rockport scale fabric

The Rockport scale fabric provides low latency and high bandwidth connectivity between compute platforms and access to storage. The architecture of the Rockport fabric ensures that all nodes are equidistant, regardless of which rack they are in. This makes design and operation much easier.

NOTE

As you review this section it may be helpful to refer to the **Example scenarios** to apply the specific concepts to see how it all ties together.

Each node in the network needs an NC1225 installed and will become a member of the Rockport fabric. This includes the compute, storage, cluster administration, and login nodes, along with the ANM and optional gateway nodes. We'll consider each of these nodes as we size the network.

Depending on the storage solution for your cluster, you may need to implement a gateway or router. If you have one of the newer software solutions that use standard server class hosts, you can easily install the NC1225 in each storage host and plug them into the same fabric as your compute servers. This delivers the best overall performance and removes a typical bottleneck in many architectures. However, if you have an existing storage solution on a separate network then you will need to set up a storage gateway or router (like an LNet Router for Lustre).

UPPER AND LOWER LEVEL SHFLS

Two SHFL variants are available to help you scale out your network:

Lower level SHFL
Rockport LS24T SHFL

This foundational piece provides interconnection for up to 24 nodes through the SHFL's green node ports. It interconnects your servers with NC1225s and can also connect to upper level SHFLs through the SHFL's three banks of black trunk ports.

Figure 2. *Rockport lower level SHFL (LS24T)*

Upper level SHFLs
Rockport US2T, US3T, and US4T SHFL

These SHFLs scale out the Rockport scale fabric beyond a single lower level SHFL and provide inter-connections to other lower level SHFLs. We'll discuss the selection and use of these SHFLs later in this document.

Figure 3. *Rockport upper level SHFLs (US2T, US3T, US4T)*

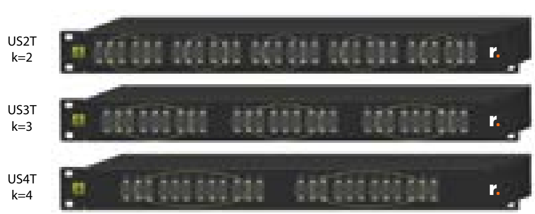

R-KEYS

When fewer than 24 nodes are connected to a lower level SHFL, green R-Keys should be connected in the unused green ports.

- **Green R-Keys** complete the topology rings, ensuring optimum topologies are created for any size of cluster.
- **Black R-Keys** create enhanced topologies, creating shortcuts between nodes in the network topology (providing a more direct path between nodes), reducing network diameter and average path length. They are inserted into unused black trunk ports on a LS24T, in cases where not all trunk ports are needed based on the target number of nodes.
- **Yellow R-Keys** (not shown) are inserted into upper level SHFL ports when there are unused ports in an alphanumeric port group, ensuring optimum topologies are created by completing the topology rings (this is typically done when an expansion is expected).

  Using the US4T as an example, if alphanumeric group 1 is cabled to three lower level SHFLs (W, X, Y port groups) and the 1Z port group is not needed based on the target number of nodes, you would install yellow R-Keys in the 1Z1, 1Z2, and 1Z3 ports. The ports with yellow R-Keys can later be replaced by a new lower level SHFL during a cluster expansion.

Figure 4. *Rockport R-Keys*

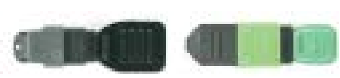

DETERMINING SHFL AND R-KEY QUANTITIES

With the total number of nodes understood we can then determine how many lower level SHFLs will be needed. This is driven by the number of nodes per rack and the number of racks. The lower level SHFLs do not need to be completely populated with connections to nodes. Use the following general guidelines to determine the number of lower level SHFLs:

- Number of nodes per rack, divided by 24 (number of ports per SHFL) and round up to a whole number
- Evenly distribute the nodes across the lower level SHFLs in each rack
- Populate the lower level SHFLs from left to right in sequence
- Install green R-Keys in unused green ports
- Avoid significant underpopulation of any single lower level SHFL. We recommend a similar number of nodes in each SHFL.

As an example, if you have 36 nodes in a rack then you can use two lower level SHFLs, each populated with 18 nodes. If that same example had four additional nodes in a second rack, then cabling those to the lower level SHFLs in the first rack would be appropriate. However, if you have eight or more nodes in the second rack, then you may consider a third lower level SHFL in this rack to avoid the operational impact of cabling across racks.

Next, we need to determine how we'll connect the lower level SHFLs together. This will depend on the total number of lower level SHFLs and the potential maximum number of nodes in the future. The lower level SHFLs each have three banks (A, B, and C) of three ports each that are used for this scaling connectivity (A1, A2, A3; B1, B2, B3; and C1, C2, C3).

You may not need all three banks to create an optimal fabric. The following table shows the best combination of upper level SHFLs based on the number of LS24Ts. For each bank (A, B, and C) you'll see the upper level SHFL type and the number of banks that will be consumed on that upper level SHFL. The upper level SHFL quantity provides the number of each.

Table 1. *Lower and upper level SHFL combinations for an optimal fabric*

Nodes	LS24T				Upper level SHFL quantity		
	LS24T qty	Bank A	Bank B	Bank C	US2T	US3T	US4T
1–24	1	Black R-Keys	Black R-Keys	Black R-Keys	0	0	0
25–48	2	US2T (1 of 5 port groups)	Black R-Keys	Black R-Keys	1	0	0
49–72	3	US3T (1 of 3 port groups)	Black R-Keys	Black R-Keys	0	1	0
73–96	4	US2T (2 of 5 port groups)	US2T (2 of 5 port groups)	Black R-Keys	1	0	0
97–120	5	US3T (1 of 3 port groups) US2T (1 of 5 port groups)	US3T (1 of 3 port groups) US2T (1 of 5 port groups)	Black R-Keys	1	1	0
121–144	6	US3T (2 of 3 port groups)	US2T (3 of 5 port groups)	Black R-Keys	1	1	0
145–168	7	US2T (2 of 5 port groups) US3T (1 of 3 port groups)	US2T (2 of 5 port groups) US3T (1 of 3 port groups)	Black R-Keys	1	1	0
169–192	8	US2T (4 of 5 port groups)	US2T (4 of 5 port groups)	US2T (4 of 5 port groups)	3	0	0
193–216	9	US3T (3 of 3 port groups)	US3T (3 of 3 port groups)	Black R-Keys	0	2	0

* Please note that the number of nodes shown in this table assumes you are fully populvating each LS24T and is intended as a guideline. The LS24T quantity dictates the number of upper level SHFLs needed.

SIZING THE AUTONOMOUS NETWORK MANAGER SERVER

The ANM is installed on a dedicated server (recommended) but can also be installed inside a Virtual Machine (VM) in the Rockport scale fabric and can be accessed through its web interface, which can be configured to any IP interface on the host. When you select the server for the ANM it is important that you consider the requirements in the following table. The ANM benefits from a dedicated local disk that is not shared with other applications.

Table 2. *Lower and upper level SHFL combinations for an optimal fabric*

Nodes	Logical cores	System memory	Disk drive requirements
Up to 24 nodes	12 CPU cores, minimum 2.0 Ghz	64 GB	100 GB /opt/rp (or other dedicated installation location) 500 GB /var/lib/docker
Up to 72 nodes	24 CPU cores, minimum 2.2 Ghz	96 GB	100 GB /opt/rp 900 GB /var/lib/docker
Up to 288 nodes	32 CPU cores, minimum 2.6 Ghz	256 GB	100 GB /opt/rp 2 TB /var/lib/docker (SSD)
Up to 648 nodes	32 CPU cores, minimum 2.6 Ghz	256 GB	100 GB /opt/rp 4.2 TB /var/lib/docker (SSD)
Up to 1536 nodes	64 CPU cores, minimum 2.9 Ghz	256 GB	100 GB /opt/rp 9.5 TB /var/lib/docker (SSD)

ANM installation instructions, more detailed sizing guidelines, and user account considerations are available in the *Installation Guide*.

SHFL CONNECTION MAPS

A SHFL connection map is used to ensure the best topology is created between the connected nodes. This map will be generated by your Cerio Solution Architect or Cerio Partner and will describe how to connect your lower and upper level SHFLs. The lower level SHFLs are numbered sequentially, with the map showing which combinations are used.

- **Bank A** refers to the first set of trunk ports on each lower level SHFL (used for dimension 4 or D4)
- **Bank B** refers to the second set of trunk ports (used for D5)
- **Bank C** refers to the third set of trunk ports (used for D6)

A map described as D4 (1, 2, 3) means that the first bank of ports from lower level SHFLs 1, 2, and 3 are connected with a US3T SHFL (because there are three in the group). The example scenarios provided in the next section will show how these are used.

Example scenarios

Earlier, we discussed the hardware requirements for the scale fabric. As we work through the design steps and decisions, we can refer to these example scenarios:

Scenario 1 – 18 node cluster 15 compute, 1 administration/login, 1 management (for ANM), 1 storage node

Scenario 2 – 40 node cluster 36 compute, 1 administration, 1 login, 1 management (for ANM), 1 storage node

Scenario 3 – 72 node cluster 67 compute, 1 administration, 1 login, 1 management (for ANM), 2 storage router nodes

Scenario 4 – 216 node cluster 208 compute, 1 administration, 2 login, 1 management (for ANM), 4 storage router nodes

SCENARIO 1

This scenario has 18 nodes in a single rack comprised of 15 compute nodes, one administration/login node, one management node to run ANM, and one storage node.

All 18 nodes will have an NC1225 installed and since we have fewer than 24 nodes, we only need one LS24T, which is shown mounted at the top of the rack in Figure 5. An active optical cable assembly (MTP/MPO-24 fiber to two QSFP-DD SR6) is attached from each NC1225 to the first available port on the LS24T. This then leaves six open node ports into which a green R-Key is installed as shown in Figure 6. Additionally, since we are only using a single LS24T, we do not need to use any of the black trunk ports which means we'll install nine black R-Keys into those ports.

Figure 5. *Scenario 1 with 18 nodes*

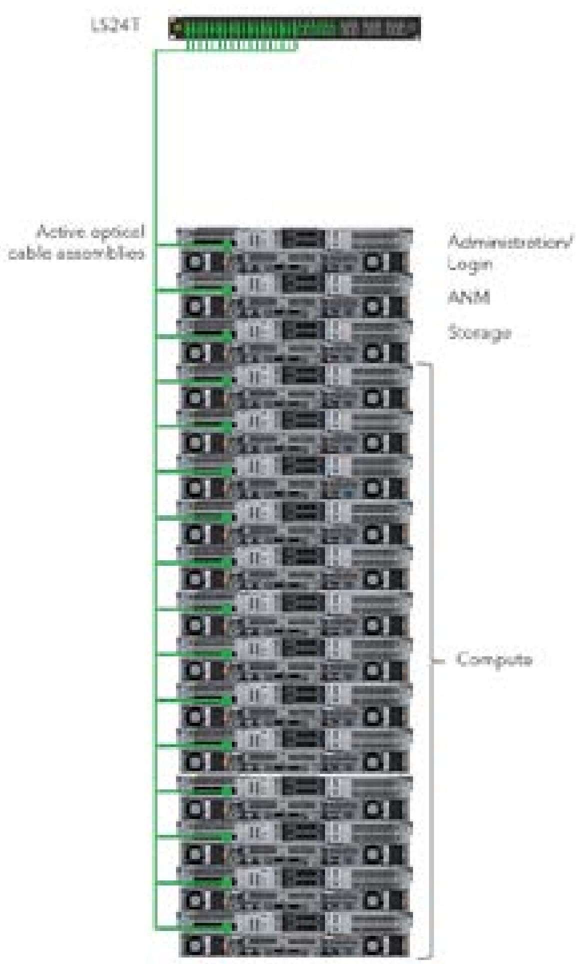

Figure 6. *Scenario 1 - SHFL detail*

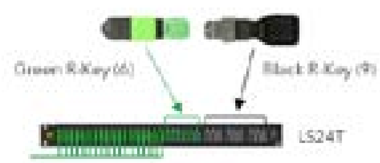

SCENARIO 2

This scenario has 40 nodes in a single rack comprised of 36 compute nodes, one administration node, one management node to run ANM, one login node, and one storage node.

All 40 nodes will have an NC1225 installed. Because we have more than 24 nodes and fewer than 48, we need two LS24Ts, which are shown mounted at the top of the rack in Figure 7. The recommended approach is to evenly split the nodes across the two LS24Ts.

An active optical cable assembly is attached from each NC1225 to the ports in the LS24Ts.

This then leaves four open node ports on each LS24T, into which green R-Keys are installed as shown in Figure 8.

At this point half the nodes are connected to each LS24T but nothing connects the two LS24Ts together. Because we have two LS24Ts, we'll use a US2T. Connecting the first three trunk ports on each LS24T to the US2T as shown in Figure 8 creates the fully connected fabric. This last step made a direct connection from each node in the first LS24T to its respective node in the second LS24T.

Additionally, since we are only using a single set of trunk ports in each LS24T we'll install black R-Keys into the six unused trunk ports of each LS24T.

Figure 7. *Scenario 2 with 40 nodes*

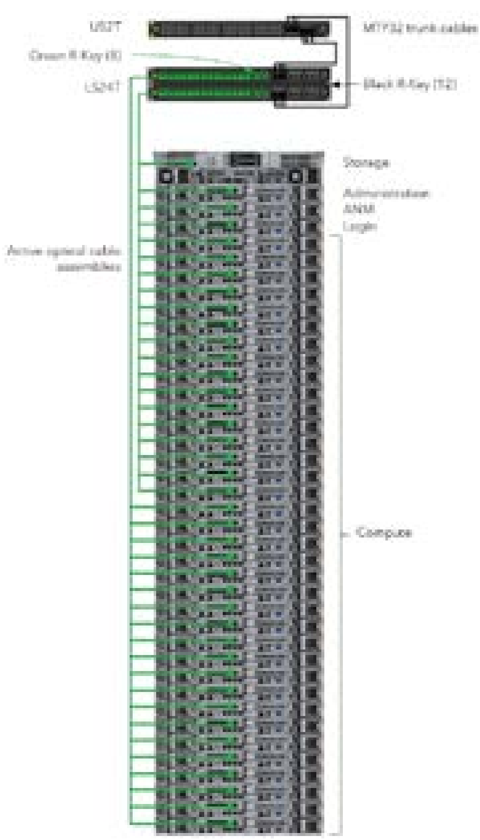

Figure 8. *Scenario 2 - SHFL detail*

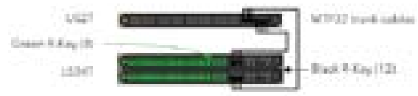

SCENARIO 3

This scenario has 72 nodes in a single rack comprised of 67 compute nodes, one administration node, one management node to run ANM, one login node, and two storage router nodes.

All 72 nodes will have an NC1225 installed. We'll need three LS24Ts, which are shown mounted at the top of the rack in Figure 9.

An active optical cable assembly is attached from each NC1225 to the ports in the LS24Ts. This fills all the ports on each lower level SHFL, meaning we do not need any green R-Keys.

At this point one third of the nodes are connected to each LS24T but nothing connects the three LS24Ts together. Because we have three LS24Ts, we'll use a US3T upper level SHFL. Connecting the first three trunk ports on each LS24T to the US3T, as shown in Figure 10 creates the fully connected fabric. This last step made a direct connection from each node in the first LS24T to its respective node in the second LS24T and third LS24T.

Since the rack is quite full, the SHFLs are mounted from a bracket hanging from the cable ladder. The product shown in Figure 10 is a Panduit Overhead Distribution Rack model #PZLRB6U.

Additionally, since we are only using a single set of trunk ports in each LS24T we'll install black R-Keys into the six unused trunk ports of each LS24T.

Figure 9. *Scenario 3 with 72 nodes*

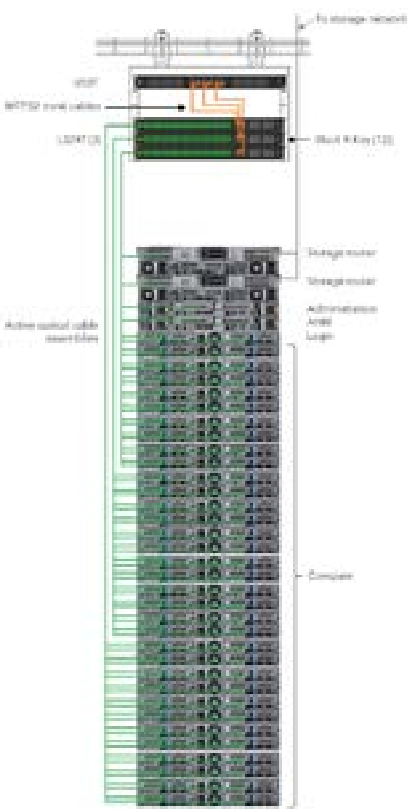

Figure 10. *Scenario 3 - SHFL detail*

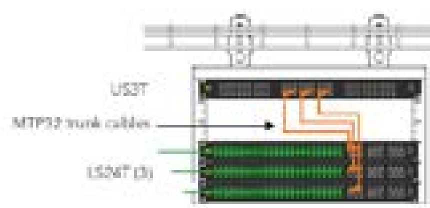

SCENARIO 4

This scenario has 216 nodes across three racks made up of 208 compute nodes, one administration node, one management node to run ANM, two login nodes, and four storage router nodes.

Note, the storage router nodes are physically located with the storage system in another rack.

All 216 nodes will have an NC1225 installed, and we'll need nine LS24Ts (three in each rack). These are shown mounted at the top of the racks in Figure 11.

Figure 11. *Scenario 4 with 216 nodes*

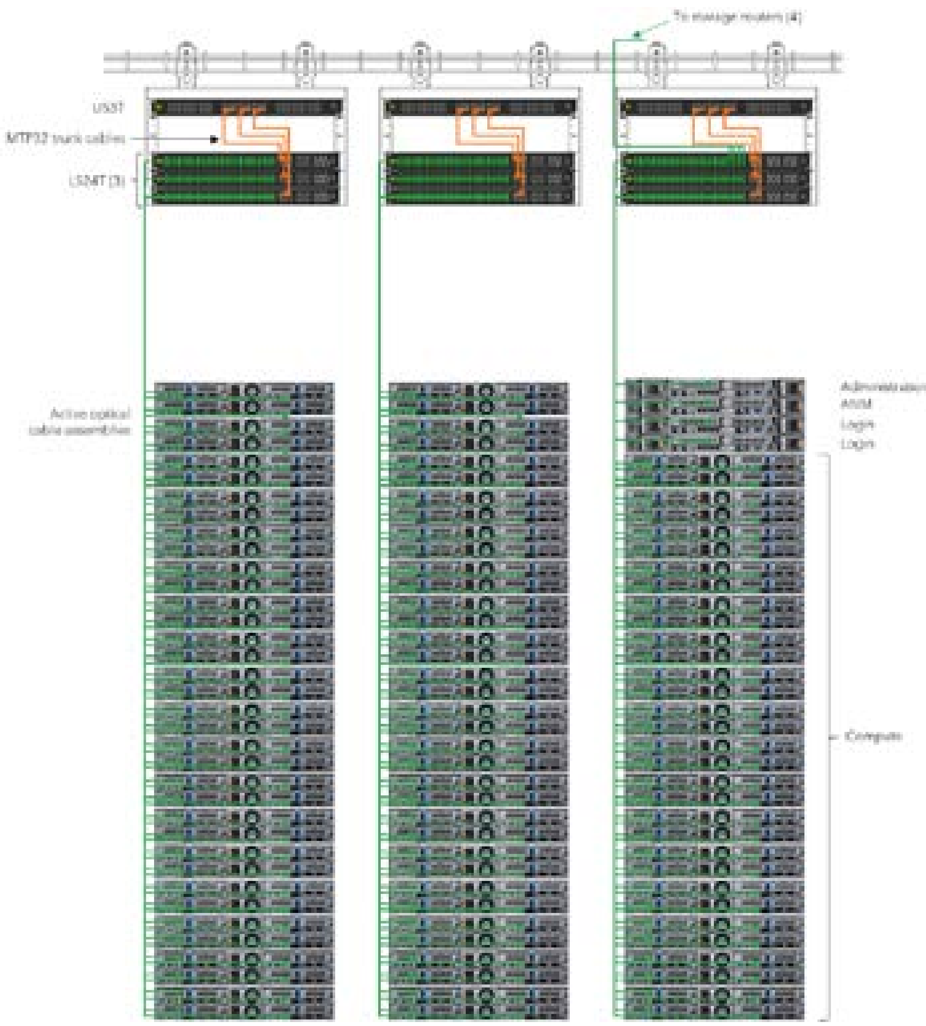

An active optical cable assembly is attached from each NC1225 to the ports in the LS24Ts. This fills all the ports on each LS24T, meaning we do not need any green R-Keys. At this point one third of the nodes in each rack are connected to each LS24T but nothing connects the three LS24Ts in each rack together. This is the same situation we had in **Scenario 3**. Because we have three LS24Ts, we'll use a US3T.

Connecting the first three trunk ports on each LS24T to the US3T creates the fully connected fabric for each group of three SHFLs. However, each of these groups are an island at this point. We must connect the three groups together. Since we have three groups, we'll again use a US3T. And since we already have a US3T with unused ports in each rack we don't have to add any more US3Ts.

Figure 12. *Scenario 4 – Connecting the first SHFL in each group*

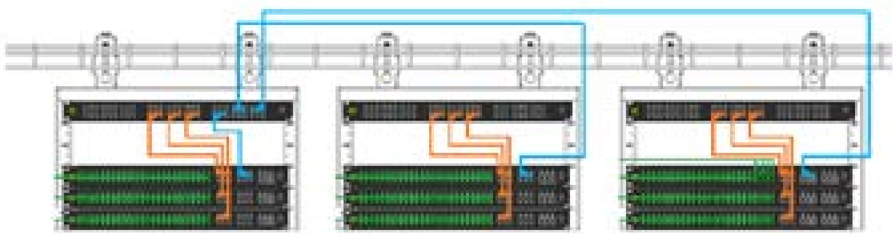

Figure 13. *Scenario 4 – Connecting the second SHFL in each group*

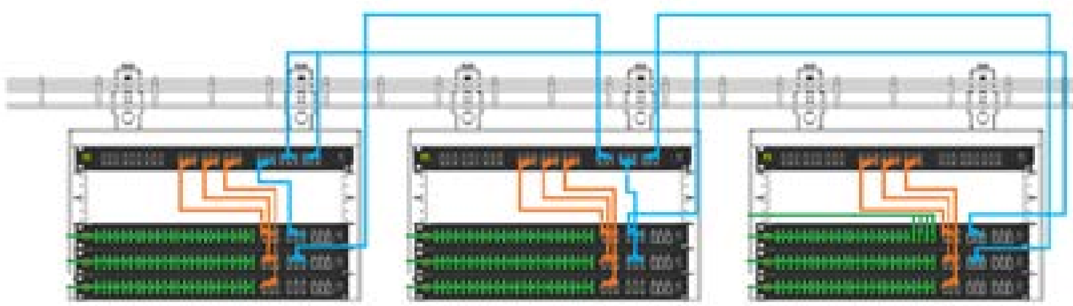

Figure 14. *Scenario 4 - Connecting the third SHFL in each group*

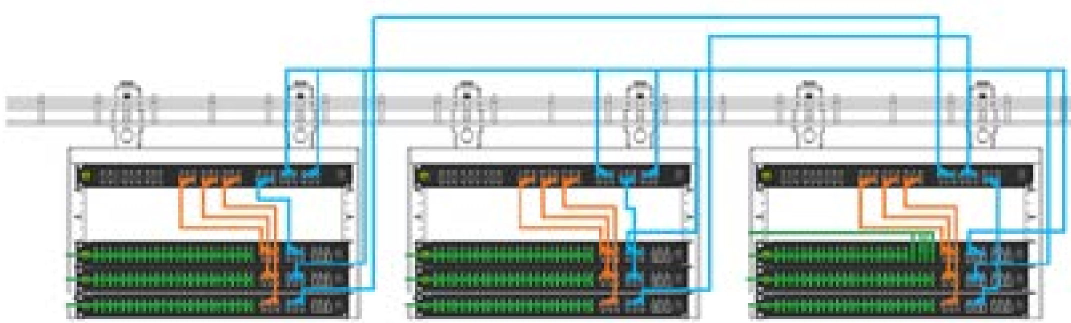

Connecting the second three trunk ports on the first LS24T in each group to the US3T, as shown in Figure 12 creates the fully connected fabric for each group of three SHFLs. We then repeat that for the second LS24T in each rack as shown in Figure 13 and for the third LS24T in each rack as shown in Figure 14. Additionally, since we are only using two sets of trunk ports in each LS24T we'll install black R-Keys into the three unused trunk ports of each LS24T.

The following tables show the BOMs for each of the scenarios described.

Table 3. *Scenario 1 BOM*

Quantity	Model	Description
18	NC1225	Rockport fabric card
18	rNOS	Rockport Network Operating System
18	ANM	Autonomous Network Manager license
1	LS24T	Lower SHFL
18	RA24G-1.5M-SR6	1.5 meter Active Optical Cable, QSFP-DD 150G SR6
1	RK24G-10PK	10 pack green R-Keys
1	RK32LB-10Pk	10 pack black R-Keys

Table 4. *Scenario 2 BOM*

Quantity	Model	Description
40	NC1225	Rockport fabric card
40	rNOS	Rockport Network Operating System
40	ANM	Autonomous Network Manager license
2	LS24T	Lower SHFL
1	US2T	Upper SHFL (two lower level SHFLs per group)
40	RA24G-x.x-SR6	Active Optical Cable, QSFP-DD 150G SR6*
2	RC32B-3.0M-3PK	3.0 meter trunk cables 3 pack
1	RK24G-10PK	10 pack green R-Keys
2	RK32LB-10Pk	10 pack black R-Keys

* Cable lengths are estimated and will vary with each deployment. 1.5 meter, 2.0 meter, and 3.0 meter cables are available.

The following tables show the BOMs for each of the scenarios described.

Table 5. *Scenario 3 BOM*

Quantity	Model	Description
72	NC1225	Rockport fabric card
72	rNOS	Rockport Network Operating System
72	ANM	Autonomous Network Manager license
3	LS24T	Lower SHFL
1	US3T	Upper SHFL (three lower level SHFLs per group)
72	RA24G-x.x-SR6	Active Optical Cable, QSFP-DD 150G SR6*
3	RC32B-3.0M-3PK	3.0 meter trunk cables 3 pack
2	RK32LB-10Pk	10 pack black R-Keys

Table 6. *Scenario 4 BOM*

Quantity	Model	Description
219	NC1225	Rockport fabric card
219	rNOS	Rockport Network Operating System
219	ANM	Autonomous Network Manager license
9	LS24T	Lower SHFL
3	US3T	Upper SHFL (three lower level SHFLs per group)
219	RA24G-x.x-SR6	Active Optical Cable, QSFP-DD 150G SR6*
18	RC32B-3.0M-3PK	3.0 meter trunk cables 3 pack
3	RK32LB-10Pk	10 pack black R-Keys

* Cable lengths are estimated and will vary with each deployment. 1.5 meter, 2.0 meter, and 3.0 meter cables are available.

Deploying and visualizing the Rockport scale fabric

From the design process you will now have a list of the necessary components (or BOM) with needed equipment, along with the ANM server requirements and the SHFL connection map. With these design details in hand, you will have what you need to deploy. The *Installation Guide* walks you through the specifics. However, these are the key steps:

1. Install the lower LS24Ts and upper SHFLs into the rack. The SHFLs do not need to be mounted inside the rack, they can be mounted above the rack or in the zero U space beside the rack.
2. Install the green R-Keys in the LS24T node ports (green) that you will not use.
3. Install the black R-Keys in the LS24T trunk ports (black) that you will not use.
4. Attach the black MTP-32 trunk cables between the lower level SHFLs and the upper level SHFLs using the connection map.
5. Install the required drivers on each of the servers and storage nodes (BEFORE installing the NC1225s). The drivers are required for the servers (compute, administration, login, ANM, etc.) and storage nodes.
6. Install the NC1225s in the servers and storage nodes in your network (if not installed previously).

   The NC1225 can be installed in a PCIe Gen3 x16 or a PCIe Gen4 x16 PCIe slot.

7. Install the ANM on the dedicated server and log in to the ANM UI. Note that it will only see itself until the other nodes come online (in the next steps).
8. Install the optical modules into the NC1225 and attach the MTP-24 end of the green cable into the first available port on the lower level SHFL.
9. As each server is powered on, you can watch each NC1225 go through the discovery and population process in the ANM UI.
10. Observe the network. The Administration Guide walks you through the details of how to use the ANM UI for day-to-day inspection and operation. Tools you can use to diagnose any issues after your installation include:
    - **Health dashboard.** Find network issues and understand the overall health status of the Rockport scale fabric.
    - **Flow Diagnostics dashboard.** See a detailed, real-time view of the traffic flowing between a source and destination node, highlighting the individual paths and links. This gives insight into your network's traffic patterns and can be a helpful tool for understanding traffic flows and investigating health issues.
    - **Path Inspector.** See all links between the source and destination node for a particular path. You can see the status of each intermediate node and port along the path, and whether alarms exist (which could contribute to issues that a node/port may be experiencing).
    - **Link Inspector.** Analyze a particular link to see detailed statistics and visualizations of bidirectional traffic between two selected ports. You can see the status of each node and port, and whether alarms exist (which could contribute to issues that the node may be experiencing). This can be a helpful tool for investigating traffic flows or health issues on a port-by-port basis, or for better understanding the utilization of a particular link.
    - **Connections that need attention.** Use the ANM's Health dashboard, Flow Diagnostics dashboard, Path Inspector, and Link Inspector to quickly find optical connections that are unseated or may need to be cleaned. When all nodes are green in the ANM UI, you can start using the network.

Further reading

There are three key documents that provide more detail and background:

Rockport Scale Fabric Technology Primer

This primer gives an in-depth overview of the Rockport scale fabric, including explanations of the key concepts of the underlying technology.

Installation Guide

This guide walks you through all aspects of a Rockport installation with a significant level of detail.

Administration Guide

This guide for administrators explains how to fully use the ANM for day-to-day inspection and operation.

We recommend that anyone deploying a Rockport scale fabric download each of these documents to use as a reference. Available from your Cerio Solution Architect or Cerio Partner.

Summary

In this reference design document, we reviewed the core components that make up the Rockport scale fabric, reviewed the necessary HPC components, discussed the architectural approach , and reviewed several examples.

As you embark on your journey with Cerio we look forward to your successful deployment. Our Partner community and Cerio Team are here to ensure your success. Please do not hesitate to contact us.

Disclaimer
Although Cerio strives to be accurate and correct to the best of our knowledge at the time of publication, the information contained herein is subject to change without notice, and is provided by Cerio on an "AS IS" basis, with all faults and errors that may appear herein. Cerio does not warrant, guarantee, accept any condition or make any representation that the functionality of the Rockport scale fabric will meet customer requirements, or that it will run uninterrupted or error free. Cerio provides no warranty against loss of data. Cerio specifically disclaims all representations and warranties to the fullest extent permitted by law, whether express or implied, or arising by law, statute, or under any other legal theory, including without limitation, warranties of merchantability, fitness for a particular purpose, satisfactory quality, and non-infringement. In no event shall Cerio be liable for any special, direct, indirect, punitive, incidental, exemplary, or consequential losses or damages whatsoever, including without limitation, loss of revenues, profits or savings howsoever arising, or those relating to interruption to business, loss of business opportunities, loss of business information or the cost of recovering such lost information, or other commercial or economic loss.

Patents and Industrial Designs
The Rockport scale fabric contains technology which is protected under various patents and industrial designs, including but not limited to: United States Patent Nos. 9,965,429, 10,303,640, 10,142,219, and 10,693,767; Canadian Patent Nos. 2,921,751, 2,951,677, 2,951,680, 2,951,684, 2,951,698, 2,951,786, and 2,939,402, and Industrial Design Registration Nos. 177010 and 177162; and European Patent Nos. 3022879 and 3087708, and Design Registration Nos. 004706943, 008034003, and 008131957. Corresponding and other patents and industrial designs are registered or pending in Australia, Brazil, Canada, China, Europe, Hong Kong, India, Israel, Japan, South Korea, Russia, and the United States. Cerio does not grant any express or implied right to you under any patents, industrial designs, or trade secret information.

Notice
The Rockport scale fabric includes proprietary software of Cerio that makes use of certain Open Source Software developed by third parties and licensed under various open source licenses. For a current listing of this open source software, associated licenses, and attributions, please refer to the "Open Source Software Notice.txt" file that is downloaded with our Autonomous Network Manager software.

ABOUT CERIO

Cerio, creating new scale economics for the AI and Cloud era, delivers an open systems platform for a more sustainable data center. Built on the foundation of a supercomputing underlay fabric, the Cerio platform provides standards-based, low-code composability services for the cost-effective and efficient management of AI/ML infrastructure.

Cerio, a trade name of Rockport Networks Inc. (DBA Cerio) is headquartered in Ottawa, Canada, with offices and projects spanning international markets, and Centers of Excellence in Europe and North America.

CONTACT US

Headquarters
515 Legget Drive, Suite 500
Ottawa, Ontario,
K2K 3G4 Canada

USA
530 Lytton Avenue, Floor 2
Palo Alto, California
94301 USA

Toll Free
1.855.771.3818

Website
cerio.io

Email
info@cerio.io

Version 3.0 October 2023