



Cornelis™ Omni-Path Express™ Fabric Design Generator

User Guide

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Cornelis Networks products described herein. You agree to grant Cornelis Networks a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All product plans and roadmaps are subject to change without notice.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cornelis Networks technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Cornelis Networks, Omni-Path, and the Cornelis Networks logo belong to Cornelis Networks, Inc. Other names and brands may be claimed as the property of others.

Copyright © 2021 Cornelis Networks, Inc.

Table of Contents

1. Introduction	5
1.1. Fabric Topologies	5
2. Generating a Design	8
2.1. Clearing a Design	8
2.2. Retrieving a Previous Design	8
3. Working with the Interface	9
3.1. Quick Config	9
3.2. Retrieve	9
3.3. Clear Config	10
3.4. Detailed Configuration Parameters	10
3.5. Configure.N	12
4. Viewing the Results	13
4.1. Show Design	13
4.1.1. Detailed Description	13
4.1.2. Brief Description	14
4.2. Show BOM	14
4.3. Configuration ID and Version	15
4.4. Download Diagram	16

Revision History

Date	Revision	Description
July 2021	4.0	Rebranded for Cornelis Networks and Omni-Path Express Fabric Design Generator name changes. Rewritten for updated content and clarity.
September 2018	3.0	Updated for Configurator 3.0.5 release (Adding Double density cables)
October 2016	2.0	Updated for Configurator release 3.0.3
July 2016	1.0	Initial release.

1. Introduction

The Cornelis Omni-Path Express Fabric Design Generator is a web-based application available for OEMs, customers, distributors, and resellers to design and compare the best fabric solutions (topologies) that meet their requirements.

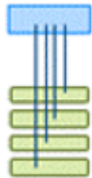
The design generator provides:

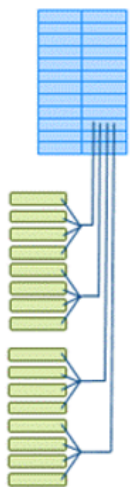
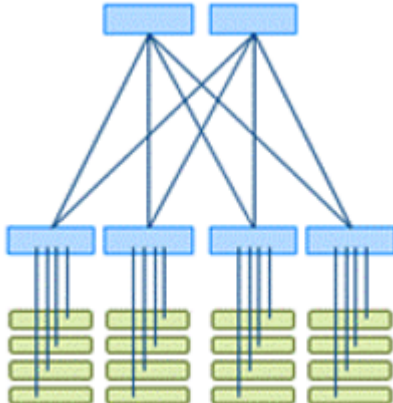
- An appropriate topology, based on user inputs
- Brief and detailed descriptions of the solution including cabling scheme and power (Watts)
- An accurate bill of materials (BOM)
- A PowerPoint diagram of the topology

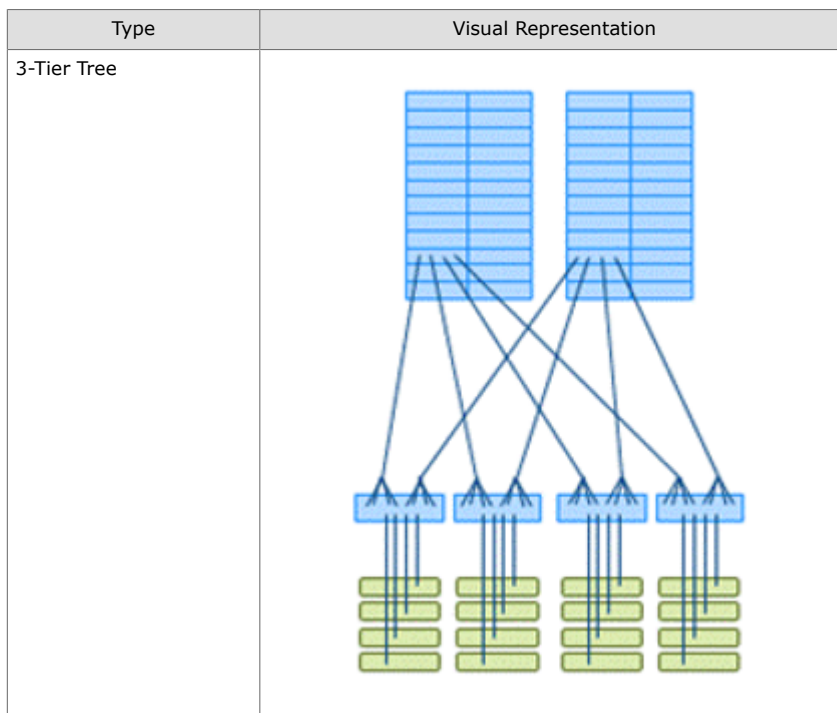
This document outlines how to use the Cornelis Omni-Path Express Fabric Design Generator. It includes example configurations as well as descriptions and diagrams of fabric topologies.

1.1. Fabric Topologies

The Fabric Design Generator will provide the solution using one of the topologies shown in the following table. You can influence the topology choice by specifying the type of switch to use.

Type	Visual Representation
Single Edge Switch	

Type	Visual Representation
Single Director	
2-Tier Tree	


NOTE

A Director chassis is simply a "fat-tree-in-box", so the 2-tier and Single Director fabric types are logically equivalent. A 2-tier fabric may be a better solution if:

- Your fabric size or the oversubscription ratio requires a significantly underpopulated Director so a Director would be relatively expensive.
- You are using a server blade system where the first tier of the fabric is built into the blade chassis.


NOTE

For 3-tier fabric trees, the Fabric Design Generator uses Director switches for the top two tiers, resulting in two tiers of physical switches. The three tiers are made from the spines, leafs and Edge switches.

2. Generating a Design

Generating a fabric design is an easy process. Just enter the number of ports for your design and the generator will provide you with multiple options.

Procedure:

Step	Task/Prompt	Action
1	Open the tool in your web browser.	The Fabric Design Generator can be found under the Cornelis Networks' Support Page .
2	Enter the number of ports for your configuration.	Click Quick Config . NOTE: Two or more configurations may be generated based on your number of ports. Click the Show Design button associated with the configuration you want to view. Review your design in the Detailed Description.
3	(Optional) To further hone your design, modify the detailed configuration parameters. Refer to Section 3.4 "Detailed Configuration Parameters" for information pertaining to each field.	For each field in the Configure.1 column, enter your parameters and click Configure.1 .
4	(Optional) To compare different configurations, enter or modify up to three additional sets of detailed configuration parameters.	Enter values in the additional configuration columns and click Configure.N where <i>N</i> is a value from 2 – 4.
5	View the bill of materials for your configuration.	Click Show BOM . NOTE: You can toggle back to the description by clicking Show Design .
6	View the fabric diagram of your configuration.	Click Download Diagram .
7	Share your configuration.	When you generate a configuration, it is automatically saved under a unique configuration ID. To share your configuration, perform either of the following actions: <ul style="list-style-type: none"> • Copy the configuration ID and send it to your distributor. • Copy your BOM results into a spreadsheet and submit it to your distributor. Refer to Section 4.2 "Show BOM" for instructions.
For information on buttons and parameters, refer to Section 3 "Working with the Interface" . To evaluate your results, refer to Section 4 "Viewing the Results" .		

2.1. Clearing a Design

To clear your configurations, click **Clear Configs**.

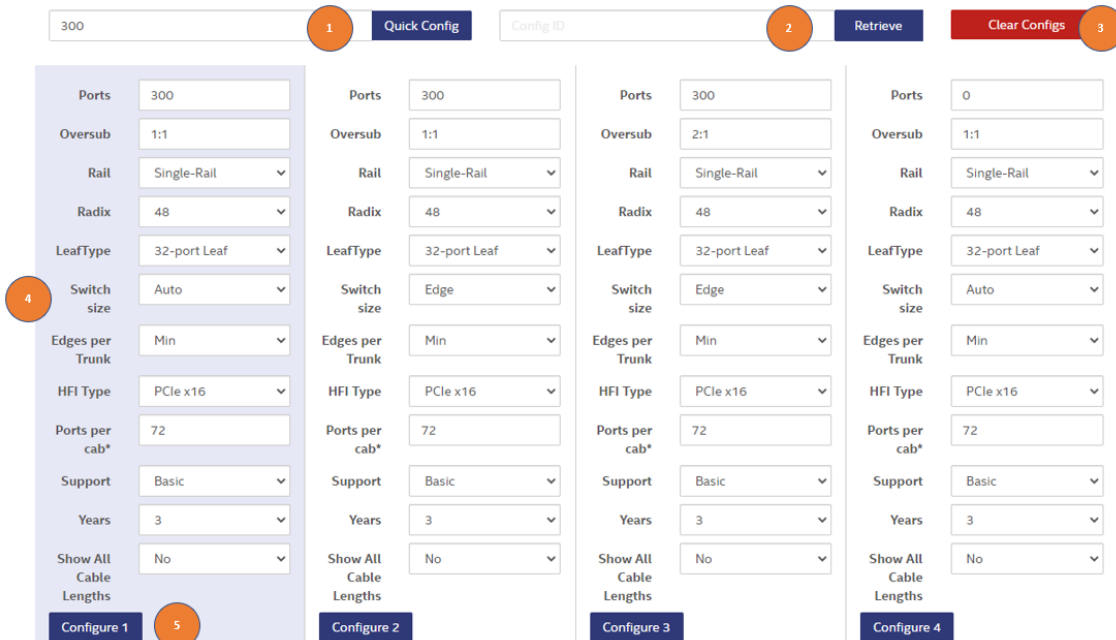
2.2. Retrieving a Previous Design

To retrieve a past configuration, enter your unique Configuration ID in the retrieve field and click **Retrieve**.

3. Working with the Interface

This section describes the buttons and parameters used to generate your fabric design.

The input area is divided into four columns, allowing you to configure four separate fabric solutions. Typically, this is so that you can compare up to four different solutions for the same opportunity.



Legend:

1. Quick Config button
2. Retrieve button
3. Clear Config button
4. Detailed configuration parameters
5. Configure.N buttons

3.1. Quick Config

The **Quick Config** button provides a number of solutions appropriate to the number of ports you have requested. These include using a Director or tree of individual switches, and optional oversubscription. Always start with Quick Config.

3.2. Retrieve

The **Retrieve** button allows you to enter an existing configuration ID to fetch a past design.

3.3. Clear Config

The **Clear Config** button resets all fields of the generator.

3.4. Detailed Configuration Parameters

The detailed configuration parameters allow you to refine your design.

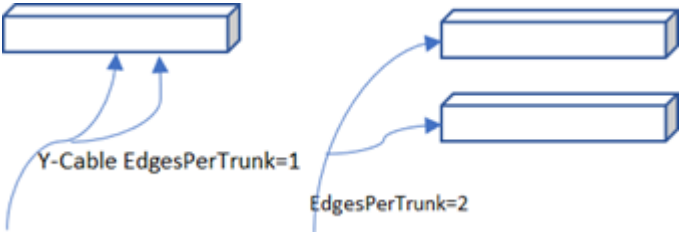


NOTE

- In general, you can set Ports and Oversub values, leaving the rest at their default values.
- If you are using Director switches, you may benefit by selecting a different LeafType value.

The following table defines each configuration parameter.

Parameter	Description
Ports	The number of end-ports required for the fabric.
OverSub	<p>The level of oversubscription required for the fabric.</p> <p>For 2-tier and 3-tier fabrics, port allocation describes how the ports on each Edge switch are divided between EndPorts and ISLports. Oversubscription can be expressed in any of these formats:</p> <ul style="list-style-type: none"> • Port Examples: 24:24, 32:16, 24:16 (where the number needs to add up to total number of ports on the Edge switches) • Ratio Examples: 1:1 (non-blocking), 2:1 • Percentage Examples: 100%, 50% <p>Interpreting Oversubscription (a:b)</p> <p>For a single Director switch, oversubscription is always interpreted as the ratio; spines will be removed from the Director switch to achieve oversubscription.</p> <p>For 2-tier and 3-tier fabric solutions, if the sum of the oversubscription value is less than 10, it is considered as the ratio, if not considered as ports. If a:b is oversubscription, and a+b is less than 10, then the ISL ports on each Edge switch is rounded up (radix (a+b)), and the ports per each Edge switch is radix-ISL ports.</p> <p>If oversubscription (a:b) is interpreted as ports, then 'b' number of cables from the Edge are connected to the core switches (ISLWidth) and 'a' number of cables from each Edge port are connected to servers (portsperEdge).</p>
Rail	<p>Options include:</p> <ul style="list-style-type: none"> • Single-Rail (default) • Dual-Rail <p>NOTE: When you select <i>Dual Rail</i>, the number of ports will be doubled.</p>
Radix	<p>The number of ports on the switch ASIC.</p> <p>For example: If Radix = X, then Edge switch has X ports.</p> <p>Default: 48</p>
LeafType	<p>Only applicable to Director switches.</p> <p>Options include:</p> <ul style="list-style-type: none"> • 32-port leaf (default) • 48-port leaf

Parameter	Description
Switch Size	<p>The size of the switch (based on number of ports) to use for the fabric.</p> <p>Options include:</p> <ul style="list-style-type: none"> • Auto (default) • Edge • 6-slot chassis • 24-slot chassis <p>When set to <i>Auto</i>, the Fabric Design Generator tool will favor a single Director switch as the solution. If you would prefer a tree of individual switches, then select <i>Edge</i>.</p>
EdgesPerTrunk	<p>This parameter is used in 3-tier fabrics built using a 48-port leaf. In a 3-tier fabric, Edge switches are connected to one or more core Director switch. The Director switches use trunk ports that contain two individual links. To connect to an Edge switch, each trunk cable splits into two tail-cables (also known as Y-Cables). The EdgesPerTrunk parameter defines how many Edge switches each trunk cable connects to.</p>  <p>Values include:</p> <ul style="list-style-type: none"> • 1: splits into the same Edge switch. This is convenient for physical cabling because each trunk cable connects to a single Edge switch allowing for short tails and therefore reducing cable density. • 2: splits into two Edge switches. This can be used for all fabrics with 48-port leaves; but it is less convenient for physical cabling because the tails of each trunk cable connect to a group of four separate Edge switches, some of which may be in different cabinets. • min (default): uses the smallest value of EdgesPerTrunk possible for this number of end-ports. The Fabric Design Generator tool first determines the minimum number of core switches required for the fabric. Next, it chooses the minimum value of EdgesPerTrunk that can be used in that topology. <p>NOTE: Cornelis Networks recommends that you always start with EdgesPerTrunk=Min. If desired, try different values of EdgesPerTrunk and compare them with the initial solution given by EdgesPerTrunk=Min.</p>
HFI Type	<p>The type of Host Fabric Interface (HFI) card required.</p> <p>Default: PCIE x16</p>
Ports per Cab	<p>The number of fabric ports in each compute-node cabinet.</p> <p>Default: 72</p> <p>Used in estimating the number of copper cables you will need in a single Director solution. Refer to the Section 4.1.1 "Detailed Description" for more details.</p> <p>NOTE: Dual rail can fit twice the number of ports per cabinet.</p>
Support	<p>The name of the service program (support level) required.</p> <p>Default: Basic</p>

Parameter	Description
Years	The number of years for which service coverage is required. Options include: 1 – 5 years Default: 3
Show All Cable Lengths	Options include: <ul style="list-style-type: none">• No (default) Shows only one type of cable in the BOM.• Yes Shows the full list of cables with different lengths supported by Cornelis in the BOM.

3.5. Configure.*N*

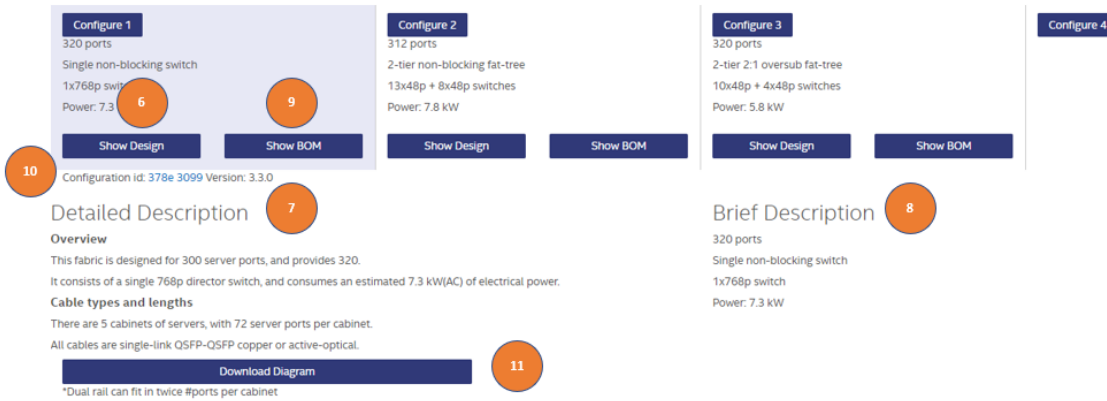
If you change any of the configuration parameters, click the **Configure.*N*** button to calculate the new solution.

where *N* is a value from 1 to 4.

4. Viewing the Results

This section describes the buttons and field used to view your results.

A brief description of the solution is shown at the bottom of each column. A detailed design description and Bill of Materials for each solution can be shown by clicking the appropriate buttons.



The screenshot displays a user interface for viewing configuration results. It features four configuration cards at the top, each with a 'Show Design' and 'Show BOM' button. Below these is a 'Detailed Description' section with an 'Overview' and 'Cable types and lengths' subsection, and a 'Brief Description' section. A 'Download Diagram' button is located at the bottom of the detailed description. Callouts 6 through 11 point to specific UI elements: 6 (Show Design button), 7 (Detailed Description header), 8 (Brief Description header), 9 (Show BOM button), 10 (Configuration ID and Version), and 11 (Download Diagram button).

Legend:

- 6. Show Design button
- 7. Detailed Description
- 8. Brief Description
- 9. Show BOM button
- 10. Configuration ID and Version
- 11. Download Diagram button

4.1. Show Design

The **Show Design** button provides the brief and detailed description of the fabric solutions built using your inputs.

If more than one configuration is available, you can click the Show Design button under each configuration to view its brief and detailed descriptions.

4.1.1. Detailed Description

The **Detailed Description** section provides in depth information about the generated configuration. In particular, it describes exactly how to interconnect the switches in 2-tier and 3-tier fabric trees. Pay close attention to the description for 3-tier fabric trees,

as the pattern of cabling to the Director leafs (described as LeafGroups) is important to the correct working of the fabric.

Detailed Description

Overview

This fabric is designed for 2000 server ports with Single-rail, and provides 2016.

It is a 3-tier non-blocking fat-tree with 2 layers of physical switches using 84x 48p edge switches and 3x 768p core switches.

It consumes an estimated 59.4 kW of electrical power.

Edges Switches

Each edge switch has 48 ports: 24 are end-ports and 24 are ISL ports.

Inter-switch cabling

The 24 cables from each edge switch are split into 3 bundles of 8.

Each of the 3 core switches takes 1 bundle of 8 cables from each of the 84 edge switches, using 672 of the 768 ports on each 768 port core switch.

Populating leaf slots

Main cores: Qty 3, each with 42x 16-port leafs. LeafGroups are formed as follows:

5 groups, where each bundle of 8 cables connects to 8 leafs, consuming 80 bundles.

1 groups, where each bundle of 8 cables connects to 2 leafs, consuming 4 bundles.

Fabric hop-count

Each server port has 23, 360 and 1632 other server ports within 1, 3 and 5 hops.

True for all servers connected to edge switches that are connected to full leafGroups

Cable types and lengths

There are 28 cabinets of servers, with 72 server ports per cabinet.

Each server cabinet contains 3 edge switches.

All cables are single-link QSFP-QSFP copper or active-optical.

4.1.2. Brief Description

The **Brief Description** section provides a shorthand description that can be used to communicate to Cornelis Networks representatives.

Brief Description

2016 ports at 24:24 (non-blocking). nCoresFilled/Factored/Packed=2.7/3/3,
edgesPerTrunk=1

84x24:24 > 3x672

Main cores: Qty 3, each with 42x 16-port leafs. LeafGroups: Grouping: 5(8L)=80B,
1(2L)=4B.

Power: 59.4 kW

4.2. Show BOM

The **Show BOM** button provides a bill of materials used to build a purchase order with latest item numbers. It also shows totals for power consumption.

If more than one configuration is available, you can click the **Show BOM** button under each configuration to view the configuration's bill of materials.

Product Code	MM	Qty	Description	Price Each(MSRP)	Total	Watts(AC) Each	kW Total
100SWDLF32Q	945777	63	Cornelis™ Omni-Path™ Director Switch Leaf Module 100 Series 32 port	0	0	189.00	11.9
100SWE48UF2	948678	84	Cornelis™ Omni-Path™ Edge Switch 100 Series 48 Port Forward 2 PSU	0	0	186.00	15.6
100SWD24B1N	945677	3	Cornelis™ Omni-Path™ Director Class Switch 100 Series 24 Slot Base 1MM	0	0	476.00	1.4
100SWDMGTSH	945776	3	Cornelis™ Omni-Path™ Director Switch Management Module 100 Series	0	0	18.00	0.1
100SWDSPINE	945778	24	Cornelis™ Omni-Path™ Director Switch Spine Module 100 Series	0	0	219.00	5.3
100SWDPS001	945780	18	Cornelis™ Omni-Path™ Director Switch Power Supply Module 100 Series	0	0		
100SWDLFFPN	945781	9	Cornelis™ Omni-Path™ Director Switch Leaf Filler Panel 100 Series	0	0		
100SWDSPFPN	945834	12	Cornelis™ Omni-Path™ Director Switch Spine Filler Panel 100 Series	0	0		
100SWD24WE1	945150	6	Cornelis™ Omni-Path™ Director Switch 24 slot Warranty Extension 1 year	0	0		
100SWE48WE1	945146	168	Cornelis™ Omni-Path™ Edge Switch 48 port Warranty Extension 1 year	0	0		
100HFA016LS	948159	2000	Cornelis™ Omni-Path™ Host Fabric Interface Adapter 100 Series 1 Port PCIe x16 Low Profile	0	0	8.00	16.0
100CQF2630	959843	2000	Cornelis™ Omni-Path™ Cable Passive Copper Cable QSFP-QSFP F 26AWG 3.0M	0	0		
100FRRL0100	952420	2016	Cornelis™ Omni-Path™ Cable Active Optical Cable QSFP-QSFP L 10.0M	0	0	4.50	9.1
Total:				0	0	Total:	59.4

2016 ports, 3-tier non-blocking fat-tree, 84x48p + 3x768p switches

To export to a spreadsheet:

Click in the Box, press Ctrl-C, then paste to Cell A1 of a spreadsheet.

*Dual rail can fit in twice #ports per cabinet



NOTE

To copy the list into a spreadsheet, click the box at the bottom of the window then press Ctrl+C to copy the information to your clipboard. Paste the information into the spreadsheet.

4.3. Configuration ID and Version

Configuration ID is a unique identifier generated for each user configuration. This ID helps users to retrieve and share the configurations with others.

Version refers to the Fabric Design Generator tool version on which the configuration was generated.

4.4. Download Diagram

The **Download Diagram** button generates a graphical representation of the cluster based on the user inputs from Fabric Design Generator tool.

2016 ports Single -rail for 2000 nodes/hosts, 3 tier, non-blocking fat-tree

