

Networking Terminology Cheat Sheet

YOUR REFERENCE SHEET FOR GOING WEB-SCALE WITH CUMULUS LINUX

With Cumulus Linux, you can build a web-scale data center with all the scalability, efficiency and automation available to the largest data center operators in the world.

You have the ability to:

- Customize your network applications
- Automate your configurations
- Choose whatever whitebox hardware you want -- based on *your* budget and *your* needs
- Build a web-scale data center like the world's largest operators at a fraction of the cost

As the new, modern way to build your network, web-scale networking gives you access to intelligent software, it's open and modular, and it allows you to automate and scale with ease.

This short guide provides you with a condensed cheat sheet of things you should consider when designing your web-scale network.


THIS DOCUMENT INCLUDES:

Data Center Switches
 Oversubscription Calculations
 Common Host to ToR Network Types
 Terminology
 CLOS Architecture
 Incremental Growth
 Optics and Cables
 TIA Standard Color Codes
 40G and 100G as Leafs
 Port Configurations


Data Center Switches

This section includes a sampling of common whitebox form-factors used in data center deployments.


TYPICAL 1RU 1G ToR SWITCH




TYPICAL 1RU 10G-T ToR SWITCH




TYPICAL 1RU 10G ToR SWITCH



TYPICAL 1RU 40G (LEAF/SPINE) SWITCH




TYPICAL 1RU 100G (LEAF/SPINE) SWITCH



Oversubscription Calculation

Oversubscription is the difference in host bandwidth (downlinks) vs. network capacity (uplinks). These calculations are shown from the point of view of a ToR Switch. A 1:1 oversubscription ratio ensures no network bottlenecks, but may result in excess capacity during non-peak times. The ideal design tries to approach 1:1 oversubscription but entirely depends on the applications and capacity needed by the administrator.


10G ToR



480 Gbps downlinks to servers / 160 Gbps uplinks to network
= 480 / 160
= 3

Which is represented as **3:1** oversubscription ratio. If every server sends at line rate, only one third of all traffic will make it into the network.

10G ToR W/6QSFP+



480 Gbps downlinks to servers / 240 Gbps uplinks to network
= 480 / 240
= 2
480 / 240 = 2


Which is represented as **2:1** oversubscription ratio. If every server sends at line rate, only one half of all traffic will make it into the network. This is better than only one third of traffic making it into the network.

THIS DOCUMENT INCLUDES:

- Data Center Switches
- Oversubscription Calculations
- Common Host to ToR Network Types
- Terminology
- CLOS Architecture
- Incremental Growth
- Optics and Cables
- TIA Standard Color Codes
- 40G and 100G as Leafs
- Port Configurations


Common Host to ToR Network Types

MLAG




MLAG is when two switches can create a single LACP bond despite coming from two independent devices. VRR (Virtual Router-Redundancy) gives the ability for a pair of switches to act as gateways simultaneously for HA (High Availability) and Active-Active mode.

REDISTRIBUTE NEIGHBOR



Redistribute neighbor daemon dynamically monitors ARP entries to redistribute these IP addresses into the fabric.

RoH (ROUTING ON THE HOST)



Routing on the host means there is a routing application (such as [Cumulus Networks Quagga](#), aptly named Cumulus Routing on the Host) either on the bare metal host or the hypervisor.

THIS DOCUMENT INCLUDES:

- Data Center Switches
- Oversubscription Calculations
- Common Host to ToR Network Types
- Terminology
- CLOS Architecture
- Incremental Growth
- Optics and Cables
- TIA Standard Color Codes
- 40G and 100G as Leafs
- Port Configurations

Common Terminology

ToR

Top of Rack switch, where servers connect to the network

PEERLINK

Link or bonded links used to connect two switches
In an MLAG pair.

LEAF

Also referred to as a ToR or Access Switch. Used typically when referring to Spine-Leaf or Clos topology.

MLAG

Multi-Chassis Link Aggregation. Ability for a pair of Switches to act redundantly in an active-active architecture and appear as a single, logical switch

EXIT-LEAF

A leaf connected to services outside the datacenter, including firewalls, load balancers and internet routers.

ECMP

Equal-cost multi-path routing. Allows load-balancing Across multiple paths.

SPINE

Also referred to as an aggregation switch, end-of-row Switch or distribution switch. Typically referred to as a Spine switch in a Spine-Leaf or Clos topology.

OOB

Out of Band Management. Indicating both the port that is used for management and the network that is dedicated to manage the infrastructure.

CLOS

A multi-stage network architecture that optimizes resource allocation for bandwidth


LAYER 3 FABRIC

Layer 3 adjacency between multiple Leafs and Spines using OSPF or BGP. Allows ECMP to take place.

THIS DOCUMENT INCLUDES:

[Data Center Switches](#)
[Oversubscription Calculations](#)
[Common Host to ToR Network Types](#)
[Terminology](#)
[CLOS Architecture](#)
[Incremental Growth](#)
[Optics and Cables](#)
[TIA Standard Color Codes](#)
[40G and 100G as Leafs](#)
[Port Configurations](#)

CLOS Architecture




THIS DOCUMENT INCLUDES:

- Data Center Switches
- Oversubscription Calculations
- Common Host to ToR Network Types
- Terminology
- CLOS Architecture
- Incremental Growth
- Optics and Cables
- TIA Standard Color Codes
- 40G and 100G as Leafs
- Port Configurations

Routing fabric is used so aggregation/spine layer can have 3 or more aggregation switches. A MLAG environment will always have a maximum redundancy of 2 and burn ports for the peerlink.

Incremental Growth

Cost-effective whitebox switches coupled with a Clos design can lead to highly redundant networks where the aggregation/spine layer is no longer restricted to 2 devices.



Optics and Cables

This section includes common optics and cabling terminology used within data centers.


THIS DOCUMENT INCLUDES:

- Data Center Switches
- Oversubscription Calculations
- Common Host to ToR Network Types
- Terminology
- CLOS Architecture
- Incremental Growth
- Optics and Cables
- TIA Standard Color Codes
- 40G and 100G as Leafs
- Port Configurations

FORM FACTORS


RJ45

8 position 8 contact (8P8C) for termination Ethernet over twisted pair. Standard termination for Cat5/ Cat6 cables.




SFP+ AND SFP28

Small form-factor pluggables. SFP+ supports 10 Gigabit Ethernet and SFP28 supports 25 Gigabit Ethernet.



QSFP+ AND QSFP28


Quad small form-factor pluggables. QSFP+ supports four channels of 10 Gigabit Ethernet QSFP28 for 40 Gigabit Ethernet links. QSFP28 supports four channels of 25 Gigabit Ethernet for 100 Gigabit Ethernet links.



TYPES OF MODULES


DAC

Directly Attached Copper. A type of cable that consists of a transceiver and cable combined together.




AOC

Active Optical Cable, a type of cable similar to a DAC where the transceiver and cable are combined except it uses multimode or single-mode fiber instead of copper.



OPTICAL TRANSCEIVER


Uses lasers to transmit and receive data over a pluggable optical fiber. The fiber can be multi-mode or single-mode depending on the type of transceiver.



TYPES OF OPTICAL TRANSCEIVERS


SR AND SR4

“Short range” transceivers use 850 nm lasers on multi-mode fiber. SR4 uses four 850 nm lasers on four parallel multi-mode fibers (four fibers in each direction).




LR AND LR4

10GBASE-LR “Long reach” transceivers use 1310 nm lasers on single-mode fiber. 40GBASE-LR4 and 100GBASE-LR4 use four different wavelengths of light centered around 1300 nm and multiplex them onto a single strand of single-mode fiber (one fiber in each direction). Some LR optics can reach 10 km in distance.



ER AND ER4

10GBASE-ER “extended range” transceivers use 1550 nm lasers on single-mode fiber. 40GBASE-ER4 and 100GBASE-ER4 use four different wavelengths of light centered around 1300 nm and multiplex them onto a single strand of single-mode fiber (one fiber in each direction). Some ER optics can reach 40 km in distance.




TYPES OF FIBER CONNECTORS

MPO

Multi-Fiber Termination Push-On. Support for both multi-mode and single-mode fiber.


MTP

Brand of MPO, sometimes used together or interchangeably.









LC

Little Connector or Local Connector. A high density fiber connector found on pluggable modules. They usually exist in pairs, one for transmit and one for receive.



TIA Standard Color Codes

Telecommunication Industry Association (TIA) is accredited by the American National Standards Institute (ANSI) to develop voluntary, consensus-based industry standards. Many data centers use cables of varying color and don't adhere to a standard.


COLOR		OPTIC
	Yellow	Single Mode Fiber (OS1 or OS2)
	Orange	Multi-Mode Fiber (OM1 or OM2)
	Aqua	Multi-Mode Fiber (OM3)
	Aqua OR	Multi-Mode Fiber (OM4)
	Erika Violet	Multi-Mode Fiber (OM4)
	Black	DAC Ethernet Cables are typically black

THIS DOCUMENT INCLUDES:

- Data Center Switches
- Oversubscription Calculations
- Common Host to ToR Network Types
- Terminology
- CLOS Architecture
- Incremental Growth
- Optics and Cables
- TIA Standard Color Codes
- 40G and 100G as Leafs
- Port Configurations

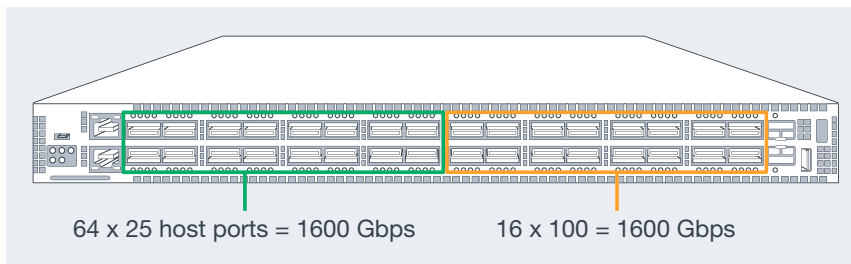
40G and 100G as Leafs

40G HIGH PERFORMANCE LEAF



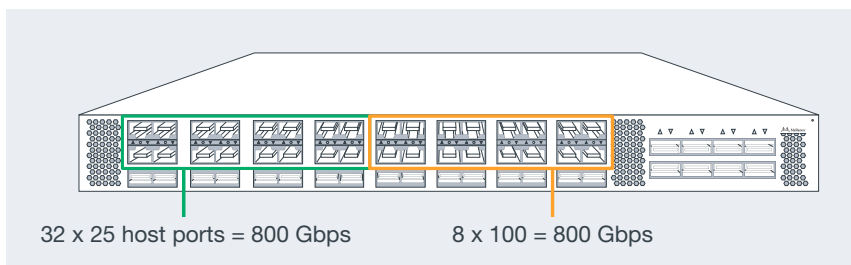
$640/640 = 1$, Which is represented at **1:1** oversubscription

100G HIGH PERFORMANCE LEAF



$1600/1600 = 1$. Which is represented at **1:1** oversubscription

100G (1/2 RU) SPECTRUM



$800/800 = 1$. Which is represented at **1:1** oversubscription

THIS DOCUMENT INCLUDES:

- Data Center Switches
- Oversubscription Calculations
- Common Host to ToR Network Types
- Terminology
- CLOS Architecture
- Incremental Growth
- Optics and Cables
- TIA Standard Color Codes
- 40G and 100G as Leafs
- Port Configurations

Port Configurations

BREAKOUT CABLE MATRIX

SFP+	Not Supported
QSFP	4 x 10 Gbps
QSFP28	2 x 50 Gbps 40 Gbps 4 x 25 Gbps 4 x 10 Gbps

BREAKOUT CABLE MAXIMUM

Broadcom Trident2:	104
Broadcom Tomahawk:	130
Mellanox Spectrum:	64*

*can fit 2 switches in 1RU for total of 128

CABLE SUPPORTABILITY

	SFP+	SFP28	QSFP	QSFP28
1GB	yes	yes*	yes	yes*
10G	yes	yes*	yes	yes*
25G	no	yes	no	yes
40G	no	no	yes	yes*
100G	no	no	no	yes

*dependent on Transceiver Compliance Codes

THIS DOCUMENT INCLUDES:

[Data Center Switches](#)
[Oversubscription Calculations](#)
[Common Host to ToR Network Types](#)
[Terminology](#)
[CLOS Architecture](#)
[Incremental Growth](#)
[Optics and Cables](#)
[TIA Standard Color Codes](#)
[40G and 100G as Leafs](#)
[Port Configurations](#)

ABOUT CUMULUS NETWORKS™

Cumulus Networks helps customers realize cost-effective, high capacity networking for modern data centers. Linux transformed the economics and innovation for data center compute, and Cumulus Linux is doing the same for the network. It radically reduces the costs and complexities of operating modern data center networks for businesses of all sizes. Cumulus Networks has received venture funding from Horowitz, Battery Ventures, Sequoia Capital, Peter Wagner and four of the original VMware founders. For more information visit cumulusnetworks.com, or follow [@cumulusnetworks](https://twitter.com/cumulusnetworks).

©2017 Cumulus Networks. All rights reserved. CUMULUS, the Cumulus Logo, CUMULUS NETWORKS, and the Rocket Turtle Logo (the "Marks") are trademarks and service marks of Cumulus Networks, Inc. in the U.S. and other countries. You are not permitted to use the Marks without the prior written consent of Cumulus Networks. The registered trademark Linux® is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis. All other marks are used under fair use or license from their respective owners.