Colocation Facility Guidelines for Deployment of Open Compute Project Racks

Authors:
Mark Dansie, Inflectiontech
Alan Keizer, AFL Hyperscale
Keith Sullivan, CS Squared
Michael J Bailey, Fidelity Investments
Stijn de Kruijf, Consultant at Royal Haskoning DHV
Robert Bunger, Schneider Electric
Jason Schafer, Google

The data center project would like to acknowledge the following people as contributors to this guidance document:

John Laban, Open Compute Project
George Cvetanovski, Hyperscalers
David Hall, Equinix
Menno Kortekaas, Circle B

1. Revision History

Date	Description
05 Oct 2017	1.0

2. Scope

This document is a reference guide for colocation data center operators and tenants who are unfamiliar with Open Compute hardware and want to understand the fundamental facility requirements to deploy this gear into their IT space.

The initial scope of the project has been focused on defining the data center sub system requirements that a European colocation facility would need to provide to accommodate an Open Compute rack. As of this revision, there are 3 types of OCP racks: Open Rack 2.0, CG Open Rack 19, and Olympus. This paper focuses on an Open Compute Rack, that when populated could weigh up to 500kg and have a maximum IT load of 6.6kW. Although, within the Open Rack design there is the capability to deploy it in all regions of the world and support a much higher IT load e.g. 36kW and up to 1400kg in weight, it was decided that to create a minimal viable product (MVP) document as quickly as possible, it would be best to restrict the checklist objectives to one that was less complex. The project team considered that if the minimum 'must-have' requirement was set at this lower level, it would allow up to 80% of the existing colo facilities in Europe to be able to accommodate an Open Rack, and therefore aid in the adoption of OCP.

3. License

This work is licensed under a Creative Commons Attribution 4.0 International License.

4. Contents

5. System Overview

Data center operators and data center tenants whose infrastructure is located in a colocation facility can take advantage of the efficiency gains made by deploying Open Compute technologies.

There are many variations on Open Compute IT gear, so it is important that the colocation operator works directly with the tenant to understand exactly what will be deployed and even guide them in the choices of OCP configurations that will fit the facility. The following chapters and the checklist in the appendices will go into detail on the various attributes of the data center subsystems and the parameters that will be required to support the current standards, Open Compute release 2.0 (aka V2), CG-OpenRack-19 and Project Olympus - that can support a base load of 500kg and an IT load of 6.6 kW.

Open Compute IT equipment looks very similar to traditional IT gear, but there are some differences that need to be taken into account. Some of these include:

- Racks arrive fully loaded: OCP IT gear is pre-installed into the rack by integrators, also known as solution providers.
- Heavy: A full OCP rack can weigh more than traditional IT equipment, so understanding facility dynamic and static floor loading limits is important.
- No centralized UPS needed: An OCP rack can come with battery back-up included, requiring the electrical feed from the facility to have no UPS upstream.
- High density: A full OCP rack can be dense, above the average for most colocation facilities, making supplying power and removing heat challenging.

This document contains guidance for an optimum sub system design, and if implemented by the colocation operator, the tenant would realize the full benefits of the various Open Rack designs.

References

- Data Center http://opencompute.org/projects/data-center/
- Rack & Power http://opencompute.org/projects/rack-and-power/
- Server http://opencompute.org/projects/server/
- Networking http://opencompute.org/projects/networking/
- Storage http://opencompute.org/projects/storage/
- Telco http://opencompute.org/projects/telco/

6. Definitions, Acronyms, Abbreviations and Units of Measurement

6.1. Definitions, acronyms, abbreviations

Access Floor		A flooring system that consists of removable and interchangeable floor tiles that are supported on adjustable pedestals and also stringers that allows access to the void under the floor (aka as a raised floor).
above finished floor (AFF)		The mounting height or vertical distance measured from the finished floor surface to the object in question.
AOC		Active optical cables are fibre cables with permanently attached transceivers. Examples are SFP+, QSFP+, CXP
Autonomy time		Is a measure of the time that a battery backup unit will support the IT gear load during a mains failure.
BBU		Battery Backup Unit
Bus Bar		The bus bars are located in the back of the rack and transmit the 12V or 48V DC power from the rack-level power sub-system to the equipment in the rack. The bars allow the equipment to plug directly into the power so the technician does not need to go to the back of the rack to disconnect power cords prior to servicing equipment.
Cat6a		Copper cabling standard that supports 10Gbe up to 90m as a permanent link (PL) + 10m of overall total patch cord length
Central Upstream UPS		The term indicates that the UPS is positioned centrally within the data center in a direction toward the utility power supply

Clos Networks		Clos networks are named after Bell Labs researcher Charles Clos, who proposed the model in 1952 as a way to overcome the performance- and cost-related challenges of electromechanical switches then used in telephone networks.
Coax		Copper network connectivity cables for short length connections 1<5m links. Examples are CX, CX-4, SFP+, CR4, CR10
Concentrated Load		The floor's capability to be able to withstand a point or static load.
Containment		Pathways through which structured cabling and power cabling is installed e.g. conduits, access floors , cable tray systems, J-Hooks and Rings
DAC		Direct attach copper
DC PDU		The DC PDU is designed to accept a nominal DC input voltage of 54V. The DC PDU box will house various power and signal connectors which will be used to power the switches and a rack monitoring device, as well as send signals to the monitoring device. There are no fuses or other secondary safety devices in this PDU.
EDA		Equipment distribution area
Fat Tree		A topology that has identical bandwidth at any bisection, each layer has the same aggregated bandwidth and each port supports the same speed as the host.
HDA		Horizontal distribution area
Inter-POD		Network connectivity between PODs of Open Racks
Intra-POD		Network connectivity within a POD of Open Racks

IT Gear		Is IT equipment installed in an Open Rack standard Equipment Bay that plugs directly into the live 12V bus bars. 'IT Gear' may also be a shelf that plugs directly in to the bus bars, and hosts multiple 'IT Trays' within the shelf. The shelf receives 12V from the Equipment Bay bus bars with one connector clip pair, and distributes the 12V to the 'IT Trays' installed within the shelf.
IT Space		Also called "White Space," this is the room which IT equipment is installed.
IT Tray		Is defined as a sub-component of the 'IT Gear' that may consist of one or more motherboards on individually removable metal trays or sleds. The mechanical and fit functions of the Open Rack standard apply only to the 'IT Gear' that plugs directly to an Open Rack bus bar system.
Leaf-spine		A two-layer network topology composed of leaf switches and spine switches. Servers and storage connect to leaf switches and leaf switches connect to spine switches. Leaf switches mesh into the spine, forming the access layer that delivers network connection points to servers/ storage. Spine switches have high port density and form the core of the architecture and allows for high performance processing of East to West data flow between servers/ storage.
LC		Latching fibre connector
MDA		Main distribution area
MM		Multimode fiber
MPO/MTP™		Multiple-Fiber Push-On/Pull-off connector for SM or MM multi-fiber ribbon. Same ferrule as MT, but easily re-connectable. Used for data center inter POD cabling and device interconnections. MTP is a brand name for a connector which inter mates with MPO.
Non-blocking fabric		Non-blocking fabric ensures that the total bandwidth of all ports that use the switch fabric

		do not exceed its capacity i.e. is the density of the ports on the switch are such that their total capacity will never be greater than that of the switch fabric. Switches operating in non-blocking mode ensure that congestion will never occur on the switch, nor will ports ever want for bandwidth between each other.
Non-blocking Clos fabric		By using the Clos mathematical theory non-blocking performance is achieved in a "switching array" (aka interconnected Ethernet fabric) without the need for n-squared ports.
OR		Open Rack
OU		OpenU is the Open Rack equivalent of the rack mount spacing found in an EIA-310 19” rack. 1 OU = 48mm vs 1 RU = 44.5mm
Out Of Band (OOB)		Is the use of a dedicated channel typically Ethernet for managing network switches, servers and storage devices.
PDU (V2)		The AC power distribution unit is designed to accept an input voltage of 3-phase WYE wiring 230/400VAC RMS voltage (4 wires + ground). The five wires will be split into two outputs via a terminal block, which will then be terminated via a Positronic connector with a custom over-molding. There will be no fuses or other secondary safety devices
POD		A group of interconnected Open Racks
Power shelf or PSU (Power Supply Unit)		An AC to DC PSU contains multiple Power Modules, typically in increments of 3 and can be operated in N+1 redundancy. The PSU has 3-phase AC input and outputs 12V or 48V DC to the rack busbars.
RPP		Remote Power Panel is the final electrical distribution panel just upstream of the racks. An RPP will typically feed a row or Pod of racks, with a 3-phase, 5 wire cable or whip to each rack.
Rolling Load		The floor’s capability to handle transportation of

		an Open Rack on casters.
SM		Singlemode fiber
Uniform Load		The load applied over the entire area of the floor in kg or lb per m^2 or ft^2.
TOR		Top of Rack (TOR) refers to a network switch which is mounted at high level in a rack.
White Glove Service		Delivery service to unload and unbox packaged racks and then transport the racks direct to the white space and place in position. The service also includes the removal of all packaging materials.
White Space		Also called “IT Space,” this is the room which IT equipment is installed.
WIP (aka WHIP)		The WIP is the electrical cable between the RPP and the power shelf within the OCP Rack. The WIP cable can be installed underneath the access floor or overhead and distributes electricity.. A WIP consists of three components 1) Receptacle -e.g. a Commando on one end, that supplies power to the input cable of the power shelf 2) Cable - that conducts electricity from the tail to the receptacle, 3) Tail – on the other end of the cable that terminates into an RPP distribution breaker.

6.2. Units of Measurement

The units of measurement used in this document are metric. Approximate conversions from metric to U.S. customary units are provided in parentheses; e.g., 100 millimetres (4 inches). Units of measurement used in this document are defined below:

C	celsius
cfm	cubic foot per minute
F	fahrenheit
ft	foot
ft^2	square foot
Gb/s	Gigabit per second
in	inch
m	metre
kg	kilogramme
kg/m^2	kilogramme per square metre
kN	kilonewton
kN/m^2	kilonewton per square metre
kPa	kilopascal
kW	kilowatt
m^2	square metre
m^3/min	cubic meter per minute
Mb	megabit
Mb/s	megabit per second
mm	millimeter
V	volt

7. Architectural

7.1. Data Center Access

When checking a colocation facility for the requirements needed to accommodate Open Racks, architectural and structural aspects need to be considered to allow a fully packaged / crated rack to be brought into the data center from the point of off-loading from the delivery vehicle, and then brought into the facility via the loading bay or dock to the goods-in area.

There are many attributes that have been considered and included in the checklist and range from a ‘must-have’ / ‘acceptable’ parameter of the delivery at road level with no step and threshold free, to a ‘must-have’ / ‘optimum’ which is a loading dock with an integral lift. The optimum design would allow packaged racks on pallets to be transported directly from inside the truck level to the data centre goods in area. The ‘must-have’/ ‘acceptable’ parameter for the delivery pathway would be 2.7m high x 1.2m wide, as this would provide sufficient height and width clearance in the doorway leading to the goods-in and unboxing locations. It is also typical for ramps to be found in colo facilities, so it is important that the gradient of any ramp in the delivery pathway is known, as a fully populated Open Rack weighing 1500kg rack would prove very difficult to move up a ramp that was steeper than a 1:12 incline.

Other ‘must-have’ attributes that can be very important to enable a smooth deployment include specifications for the delivery pathway within the data center, such as height and width of door openings in corridors, and the maximum weight a lift can carry.

Open Rack Packaging

	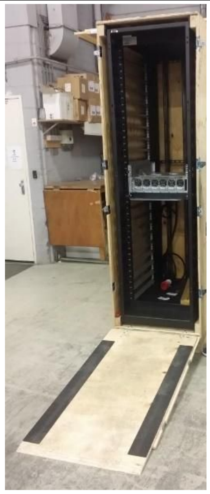	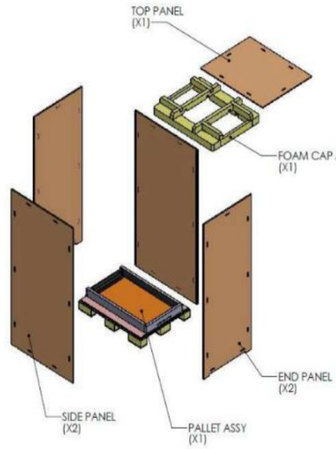	
Rack upon arrival	Crate with integral ramp	Exploded view of a packing typical crate	Loaded rack in a crate with front / rear sides removed

Delivery Pathway

	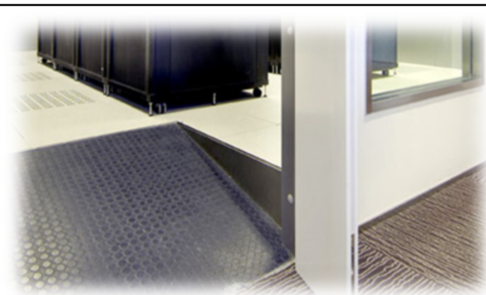
Acceptable Ramp Gradient	Unsuitable Ramp Gradient

	
Spreader plates can be used to increase rolling load capacity	Example of improper floor strengthening

7.2. IT Space (White Space)

In the checklist, a number of structural attributes for a data center white space have been considered, with many classed as 'must-have'. Open Racks are heavy in nature and many of the traditional colos built even as recently as 10 years ago were not designed to accommodate Pods of 24 racks, with each rack weighing between 500kg to 1500kg, so a 'must-have' / 'acceptable' parameter for the access floor uniformed load to support a 500kg rack would be 732 kg/m^2 (150 lb/ft^2)(7.17 kn/m^2).

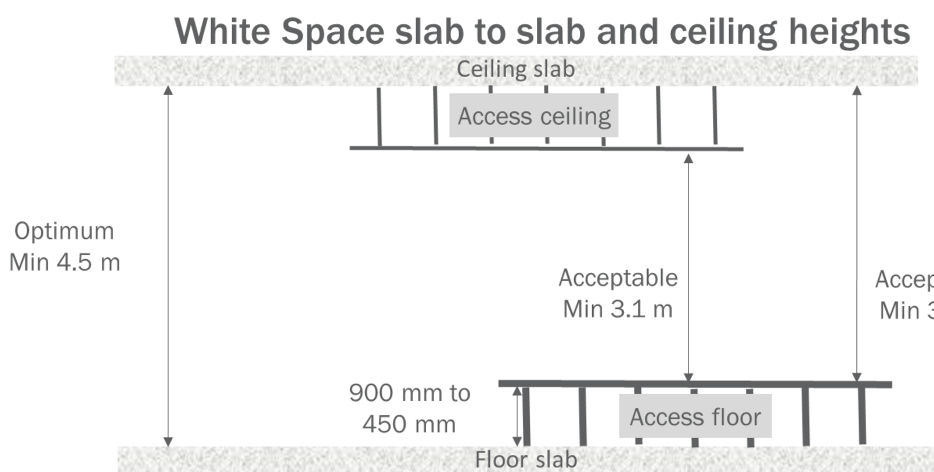

8. Cooling

Cooling an Open Rack is very similar to a rack of traditional IT gear. It has front to back airflow and the density can vary from low to high (~4 to greater than 20kW) depending on the configuration of equipment. The amount of airflow per kW of load can vary based on firmware and the design delta temperature across the server, so it's important to understand where in the spectrum your IT gear falls.

As a best practice, containment is recommended for any density to improve energy efficiency. Hot aisle containment is considered to be the optimum design, since it can be supported on hard floor facilities. Also one of the many advantages of the Open Rack design is that all servicing and cabling of the equipment in the rack can be carried out at the front, so if the racks are contained in a hot aisle then maintenance personnel will need never enter that space, which is normally very uncomfortable to work in. Therefore it has been considered as a 'nice-to-have'/

‘optimum’ arrangement to have a hot aisle containment system. The ‘must-have’ attributes in this section of the checklist include either hot aisle or cold aisle containment, front to back airflow and inlet temperature, and humidity within the Ashrae recommended limits.

IT space air containment examples

	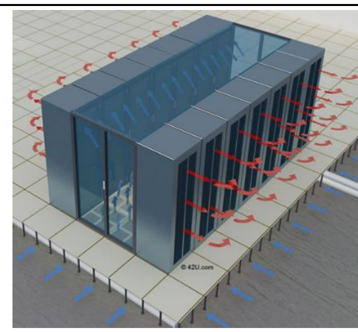
Optimum Design	Acceptable Design

9. Electrical Systems

The IT gear within an Open or CG19 Rack is powered by one or two rack mounted power shelves or PSUs[1], containing AC to DC rectifiers, which distribute 12V or 48V via busbars in the back of the rack to the equipment. The OCP rack can also contain Lithium Ion batteries on a shelf (BBU) that would act as the battery backup and therefore providing a benefit for a colo to not have to provide a centralised upstream UPS. For a colocation facility in Europe to be able to accommodate an Open Rack that has an IT load of 6.6kW, a ‘must-have’ / ‘acceptable’ requirement would be to provide a rack supply, fed by a central upstream UPS with a capacity of 3-phase 16 amp, with a receptacle compatible with IEC60309-2 5 wire. The ‘nice-to-have’ attribute, which has been categorised as ‘optimum’ within the checklist, as it provides an opportunity to be more energy efficiency and resilient, would be for the facility to provide a supply to the rack that was not from the central upstream UPS, but from the UPS input distribution board. Considerations for a

[1] Project Olympus does not use centralized rack PSUs but takes 3-phase AC directly to the server via specialized blind mate rack PDU.

colo and tenant would be to understand the generator start up time if the racks were reliant on the battery backup unit (BBU) of the power shelf to be the UPS, so as to ensure that there was sufficient autonomy time to keep IT gear functioning before the genset comes online.

3-phase circuits will be most commonly used to supply Open Racks. A 32A 3-phase supply will cover most scenarios, with 16A 3-phase as the ‘must have’ parameter for low power applications requiring 6.6kW. IEC 309 connectors will be used to connect to the power shelf. Consult with an OCP solution provider to understand circuit and plug-type required. Many OCP V2 configurations will have equipment that is not powered from the 12V or 48VDC busbars, and will need AC power, such as top of rack switches.

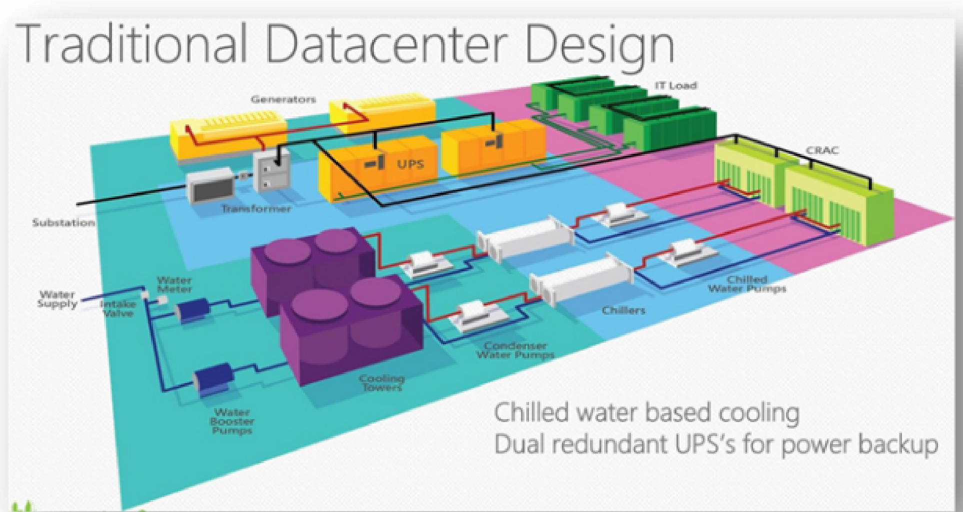

Acceptable Design

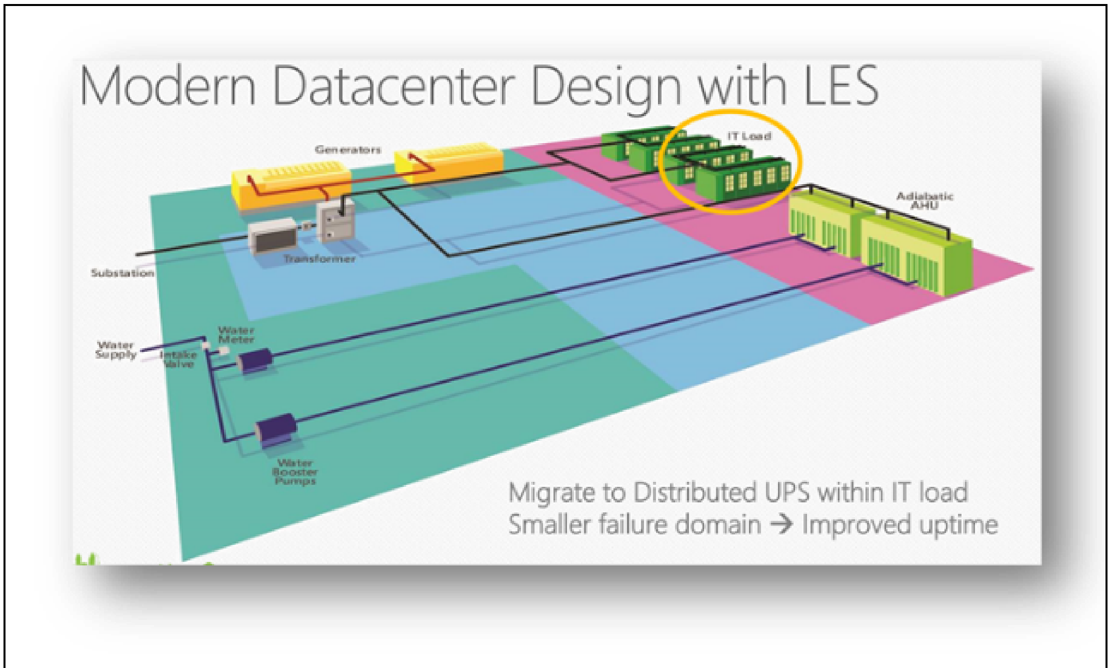

Optimum Design

10. Telecommunication Cabling, Infrastructure, Pathways and Spaces

There are a number of considerations to take into account, for example, can power and network connections be moved from the rear to the front of the rack. This is more complicated than it seems at first. Depending on how the communication and power cables run, it may be necessary to enable the change from standard EIA 19" rack to an Open Rack by breaking one cable at a time, re-routing, re-connecting and testing. In addition, OCP racks are rolled into place fully configured, so the facility infrastructure (power, networking, containment) needs to be set up in a way to allow for easy installation and removal of full racks.

Can network and power connections be routed from the access floor to the top of the rack as in many traditionally built data centers power and communication cabling is routed underfloor? The 'must-have'/ 'acceptable' arrangement for routing network cabling into an Open Rack would be either top or bottom entry and to the front of rack. The V2 rack has an exit hole in the bottom of the rack 25mm x 15mm minimum in size under each cable zone in the front of the rack to allow for a data cable under the floor to pass to the cable zone.

The network design of the data center design is very much specific to the needs of the tenant's use case. Attributes to be considered by the tenant include maximum link distance between Spine & Leaf network switches, transmission speeds of TOR switches, media type for TOR to Leaf and Leaf to Spine connectivity.

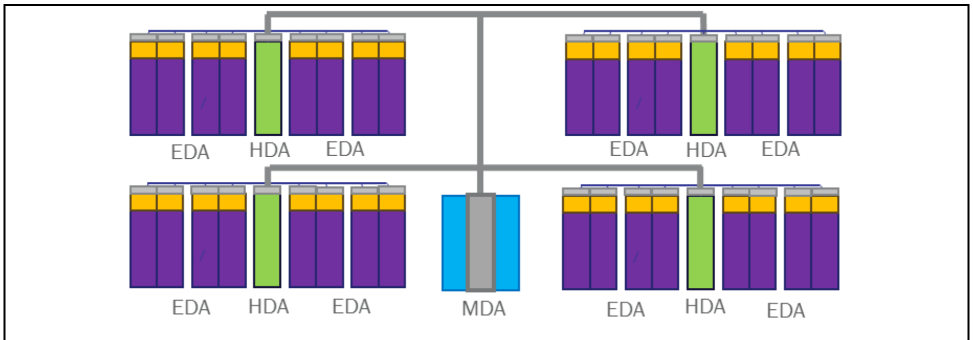

Consideration for Space Design

Optimum Pathway Design

Appendices

A. Checklist for an OCP Rack

a. Introduction

Within the checklist the attributes of each data center sub system have been assessed and listed in rows, below one of the classification headings of 'must-have', 'nice-to-have' or 'considerations', and the parameters of each attribute have then been inserted into one of two columns, with the headings of 'acceptable' or 'optimum'. The 'must-have'/ 'acceptable' attributes and their corresponding parameters have been considered to be the minimum requirement needed to be provided by an European colo facility to accommodate an OCP Rack, which weighs a maximum of 500kg when populated, and an IT load of 6.6kW. Note, OCP rack configurations can easily exceed this weight and power density, so this should be considered the minimum acceptable specification. The 'nice-to-have' attributes are viewed as not essential for a deployment, but could be beneficial based on a particular scenario. The attributes under the classification heading of 'considerations are those which are usually tenant specific requirements. There is also guidance information within the checklist for an attribute's parameter to be considered as optimum, and if implemented by the colocation facility or tenant would enable the full benefits of the OCP Rack design to be achieved.

b. Checklist

The printable checklist can be found on the OCP Datacenter wiki:

http://opencompute.org/wiki/Data_Center

B. Open Compute Rack Types

There are several variations on OCP IT gear, so it is important that the operator of colo facility works directly with the tenant to understand exactly what will be deployed and even guide the choices of OCP configurations to fit the facility. Below is a table showing the three main Open Compute rack types and their fundamental specifications.

Attribute \ Rack Type	Open Rack v2.0	CG-OpenRack-19	Project Olympus
outside width	600mm typical, but can vary	600mm	Any Standard EIA-310 rack
depth	Standard: ~1048mm Shallow: ~ 762mm	1200mm	Any Standard EIA-310 rack
height	unspecified, but facebook uses 2210mm	variable	Any Standard EIA-310 rack
weight	variable 500 - 1400kg	variable 500 - 1500kg	TBD
mounting rail spacing	21”	19” Standard EIA-310	19” Standard EIA-310
U Spacing	48mm OpenU (OU)	44.45mm (1.752”) Standard EIA-310	44.45mm (1.752”) Standard EIA-310
Required access	primarily front only	primarily front only	primarily front only
PSU architecture	3 phase AC rack PSU to 12V or 48 VDC busbar distribution	3 phase AC rack PSU to 12VDC busbar distribution	3 phase PSU internal to server
Battery Backup	optional in-rack Li-ion	optional in-rack Li-ion	optional in-server Li-ion
Power feed to rack	Typically 3 phase AC 230/ 400 VAC	3 phase AC 90 - 264VAC	3 phase AC 230/ 400 VAC
Airflow	front to back	front to back	front to back

Image	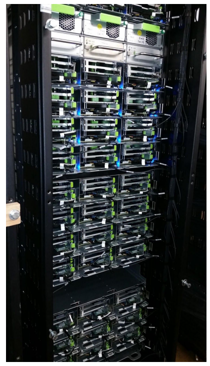		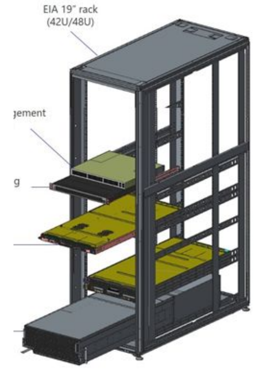

C. OCP Principles

Open Compute Project embraces 4 principles: Efficiency, Openness, Scale, & Impact.   All contributions made to the OCP foundation are evaluated against these principles and must demonstrate that the product contribution meet some or all these.   The four principles are:

Efficiency - A key OCP tenant is efficient design. Aspects considered include (but aren't limited to) power delivery and conversion, thermal efficiency, platform performance (per-W for example), reducing overall infrastructure costs, reducing code weight, reducing latencies and more.

Scale - OCP contributions must be scale-out ready. This means that the technology is designed with the right supporting features to allow for its maintenance in large scale deployments. This includes physical maintenance, remote management, upgradability, error reporting and appropriate documentation. Management tools should strive to adhere to the guideline provided by the OCP Hardware Management Project. Documentation should enable adopters towards a successful deployment, providing guidance on equipment installation, turn on and configuration, as well as physical and remote service.

Openness- OCP encourages as much open source contribution as possible. Whether fully open source or not, a contribution should strive to comply with a set of already existing open interfaces, at the very least provide one. Providing a solution compatible with already existing OCP contributions is one way to implement existing (open) interfaces.

Impact - New OCP contributions must create meaningful positive impact within the

OCP ecosystem. This can be attained by introducing efficiency gains, introducing new technologies and products that are valuable for scale out computing, creating a multiplier effect by building on top of already existing OCP solutions, and enabling a more robust supply chain by contributing alternative compatible solutions.

D. Cable Plant Overview

How should you lay out your cable plant? There are potentially 100 different ways to do this but let's look at three that have been used most commonly: Home Runs, End of Row, Centralized Distribution.

1. Home Runs: This is a practice of running jumper cables from host to switch. Home Runs are the most cost efficient cable plant that you can install, but it has many drawbacks. This option is not scalable, troubleshooting is problematic, and moves, adds, & changes are risky. This would only be recommended for very small installations, such as 1 to 2 racks.

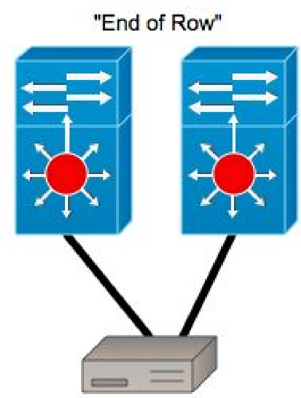

2. End of Row: Commonly used in U shaped network topology with chassis style switching. Cabling would leave cabinet from either top or bottom traversing the cable tray to the end of its row. This cabling would be terminated into a patch field adjacent to the switch. Patch cords would be used to wire up hosts from patch panels to switch to complete the link.Fiber cabling would be used to link the end of row switches back to the core and then onto the WAN. This option is good for medium to large Data Centers offering them the ability to be strategic and flexible to change. The downside to this design is that you are locked into your cabinets with this cabling. When technology updates it may require a rebuild of the cabling at the each rack or the whole cable plant.

3. Centralized Distribution: This could be the most expensive option but it is the most flexible. This layout ensures that you are ready to receive whatever the data center has to throw at you. Cabling will be installed above the racks attached to the cable tray. The cabling would traverse the cable tray back to the centralized patch field. Ideally, this patch field would be laid out to make a map of the room. The patch field would also give you the

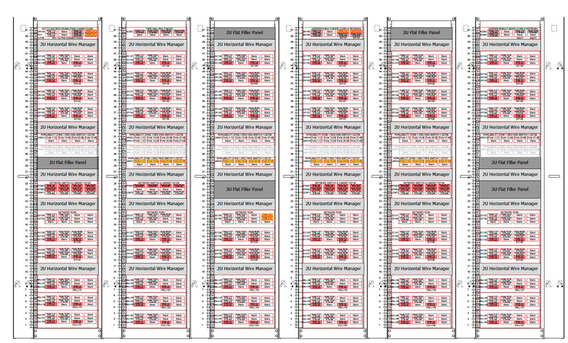

ability to jump from rack to rack, row to row, and room to room. Commonly used with top of rack but will support U network topology as well.

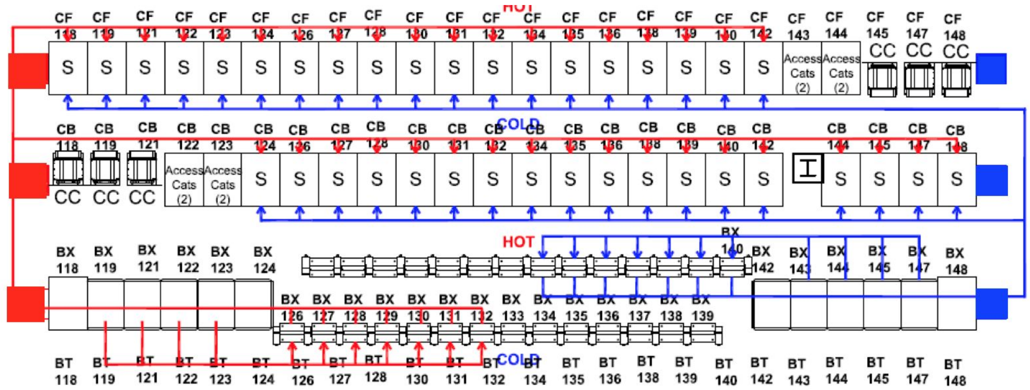

E. Cable Plant Reference Architecture

Introduction

A reference architecture for Open Racks which has been designed to have basic functionality, with a network topology that is a two or three layer Clos Fat Tree network which is scalable between 100 to over 100,000 servers. Other notable aspects of the reference architecture are that overhead pathways are used to allow easy full rack exchange and front and top data cable and service access allows for hot aisle containment to be used for cooling.

Leaf and Spine Fabric

2K Array” Example

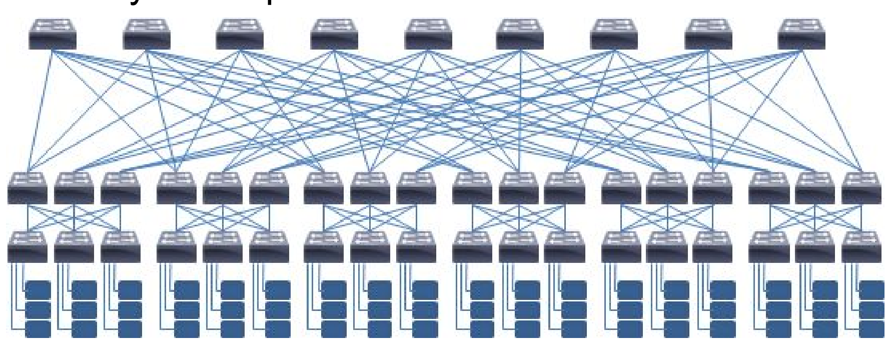

- Many leaf and spine variations are possible. One example is the “K Array” shown here for reference.
- “Fat Tree” characterised by same bandwidth at each layer.
- Folded mesh allows extensive scaling with equal delay connections.

- Non-blocking fabric assuming uniform packet distribution.
- Single switch type throughout.
- Each port supports same speed as end server.
- K-port switch supports $K^3/4$ servers.
- Example above; K=6 with 6 pods holding 54 servers, 45 total switches and 270 total ports.
- Smaller fabrics may have only two layers.

K Array Metrics

	Per Pod			Totals						
K	Servers	Access	Aggre-gation	Servers	Access Switches	Aggregation Switches	Core Switches	Total Switches	Total Ports	Total Channels
4	4	2	2	16	8	8	4	20	80	40
6	9	3	3	54	18	18	9	45	270	135
8	16	4	4	128	32	32	16	80	640	320
12	36	6	6	432	72	72	36	180	2,160	1,080
16	64	8	8	1,024	128	128	64	320	5,120	2,560
24	144	12	12	3,456	288	288	144	720	17,280	8,640
32	256	16	16	8,192	512	512	256	1,280	40,960	20,480
48	576	24	24	27,648	1,152	1,152	576	2,880	138,240	69,120
64	1,024	32	32	65,536	2,048	2,048	1,024	5,120	327,680	163,840
96	2,304	48	48	221,184	4,608	4,608	2,304	11,520	1,105,920	552,960
128	4,096	64	64	524,288	8,192	8,192	4,096	20,480	2,621,440	1,310,720

Leaf/Spine Physical Design 1024 Servers Example

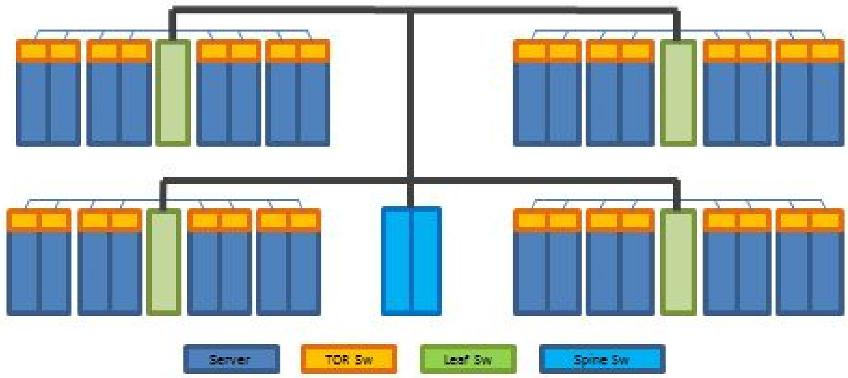

Example

- 16 port switches (K=16)
- 16 pods with 64 servers each – 1024 total servers
- 128 access switches, 128 aggregation switches and 64 core switches
- 32 servers + 4 access switches per rack => 2 racks per pod [32 intra rack and 32 external channels per rack]
- 4 clusters of 4 pods; each cluster having 32 aggregation switches in a single rack [512 channels per cluster; 256 to access switches and 256 to core switches]
- 64 core switches in two racks [1024 channels]

Structured Cabling View 1024 Server Leaf/Spine Example

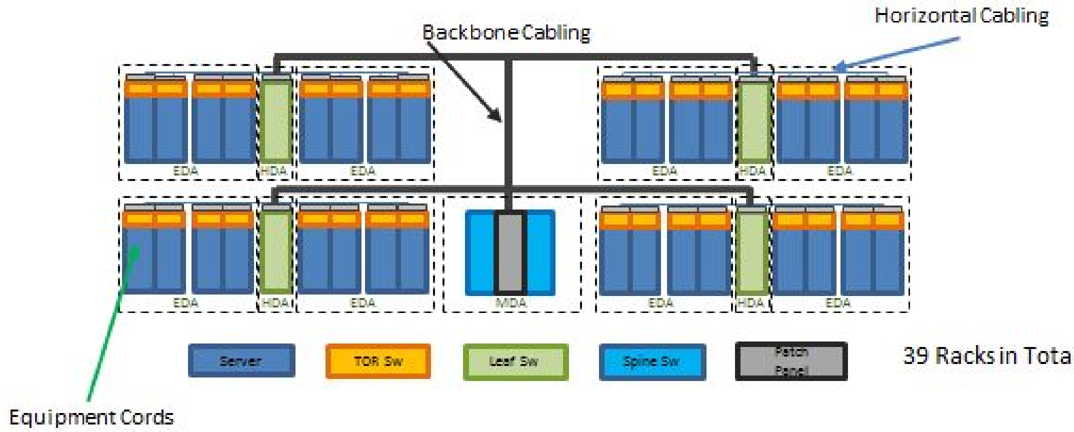

Equipment Cords *10G and 25G* *DAC, Cat6A, MM*	**5120**	**EDA Patch Panels** *Duplex LC, LC/MPO Modular*	**32 x 32 Ports**
Horizontal Cables *AOC, MM, SM*	**32 x 32 Channels**	**HDA Patch Panels** *Duplex LC, LC/MPO Modular*	**4 x 512 Ports**
Backbone Cables *SM, MM (smaller configurations)*	**4 x 256 Channels**	**MDA Patch Panels** *Duplex LC, LC/MPO Modular*	**1 x 1024 Ports**