Open Compute Project

Colocation Facility Guidelines for Deployment of Open Racks

Authors:
Mark Dansie, Open Compute Project Foundation
Stijn De Kruijf, Royal Haskoning DHV
Sean Halpin, Dataplex Group
Gerard Thibault, Kao Data
Brevan Reyher, Rackspace
Rob Coyle, PCX Corporation

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

Gaby Mindreci, Positronic Europe
Robert Bunger, Schneider Electric
Stephan Harren, Maincubes
Lucas Cardholm, Coromatic
John Laban, Open Compute Project Foundation
Jean-Marie Verdun, Hewlett Packard Enterprise

1. Revision History

Date	Revision	Description
05 Oct 2017	1.0	First Final
15 Aug 2018	1.1	Working Draft
22 Oct 2018	2.0	Revised Final for Publication
27 Apr 2020	3.0	Working Draft
7 Dec 2020	3.1	Final for Comment
15 Dec 2020	4.0	Final for Publication

2. Scope

The initial scope of the sub-project that started in 2016, was focused on defining the data center sub-system requirements that a colocation facility would need to provide to accommodate an Open Compute Project Open Rack.

This document contains guidance for an optimum sub-system design, and if implemented by the Colocation Solution Provider (CSP), the OCP user would realize the full benefits of the various Open Rack designs.

This document is also a reference guide for those who are unfamiliar with Open Compute hardware and want to understand the fundamental facility requirements to deploy this gear into their IT technical space.

As of this revision, there are two main types of Open Racks:

- Open Rack v2.0
- Open Rack v1.0

The project team identified the minimum must-have attributes and their parameters, defined as an 'Acceptable' level. This level assumes that up to 80% of the existing colo facilities are able to accommodate Open Rack deployments.

Also a set of more future proof, mature level of parameters for attributes were defined as the 'Optimum' level.

As reference for the levels set for the attributes:

- 'Acceptable' Open Rack, that when populated could weigh up to 500kg and have a maximum IT load of 6.6kW.
- 'Optimum' Open Rack, that when populated could weigh up to 1500kg and have a maximum IT load of 36kW

As the chapters of the Guidelines are developed, to make the cross referencing of the OCP Ready Requirements easier to follow within these guidelines, an OCP Ready Requirements Matrix reference (matrix ref.) will be quoted in the text. The OCP Ready Requirements matrix itself can be found at Appendix F.

Please suggest change proposals or content additions in the following document:
'Change Proposal - OCP Ready'
https://docs.google.com/forms/d/e/1FAIpQLSfHvryfaV8sYsAXQ4MJ1z1B7Goi7OLXg0lCywVOL2o8OTSqQA/viewform

3. License

4. Contents

5. System Overview

OCP users whose infrastructure is located in a colocation facility can take advantage of the efficiency gains made by deploying Open Compute technologies.

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 populated: OCP IT gear is pre-installed into the rack by integrators, also known as OCP Solution Providers (SPs) that have Level 11-12 manufacturing capability https://www.amax.com/blog/?p=668
- 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 (BBU) included, requiring the electrical feed from the facility to have no centralized UPS upstream.
- High density: A full OCP rack can be dense, above the average for most colocation data centers, making supplying power and removing heat challenging.

There are many variations on Open Compute IT gear, so it is important that the Colo Solution Provider (CSP) works directly with the OCP user to understand exactly what will be deployed, and even guide them in the choices of OCP hardware configurations that will be suitable for the facility.

This document contains guidance for the OCP Ready requirements for an optimum sub-system design, and if implemented by the CSP, the OCP user would realize the full benefits of the various Open Rack designs.

The following chapters and the OCP Ready requirements matrix in the appendices, will go into detail on the various attributes of the data center sub-systems and the parameters that will be required to support the current requirements.

In this document the OCP Ready requirements have been listed for:

- Data Centre Access
- IT Technical Space (White Space)
- Electrical systems
- Cooling system
- Telecommunication cabling, infrastructure pathways and spaces

6. Data Center Access (Matrix Ref. 1)

When checking a colocation facility for the requirements needed to accommodate Open Racks, vehicular access, 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.

6.1 Vehicular Access

In order to facilitate not only the daily traffic to the building, but also delivery of racks and equipment to the site, the turning radii of commercial vehicles must be considered when allowing for traffic usage. Vehicular access to and from the site for the purpose of loading and unloading racks is a key attribute to any logistical plan for the Data Hall fitout.

6.2 Delivery Pathway

There are many ‘must have’ attributes that have been considered and included in the OCP Ready requirements matrix and range from an ‘acceptable’ parameter of the delivery at ‘road level with threshold’ with mitigation e.g. use of a temporary ramp, to an ‘optimum’ which is a loading dock with an integral lift (matrix ref. 1.01). 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 ‘acceptable’ parameter for the delivery pathway would be 2.3m (90in) H x ≥ .9m (36in) W with threshold mitigation (matrix ref. 1.02), 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 (3,300 lbs) would prove very difficult to move up a ramp that was steeper than a 1:12 incline (matrix ref. 2.01).

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 (matrix ref. 1.03), and the maximum weight a goods lift can carry (matrix ref. 3.01).

The ‘acceptable’ and ‘optimum’ parameters for the rolling, uniform and concentrated loads have been derived from the performance specifications listed in ANSI/BICSI 002-2019 and The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992.

6.3 Open Rack Packaging

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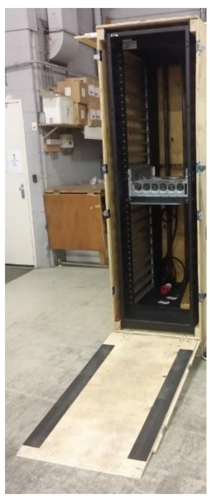</td>
<td>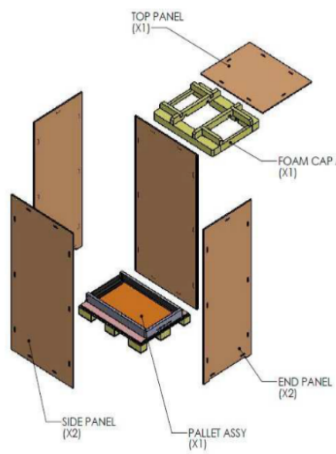
</td>
<td></td>
</tr>
<tr>
<td>Rack upon arrival</td>
<td>Crate with integral ramp</td>
<td>Exploded view of a packing typical crate</td>
<td>Loaded rack in a crate with front / rear sides removed</td>
</tr>
</table>

6.4 Delivery Pathway Examples

	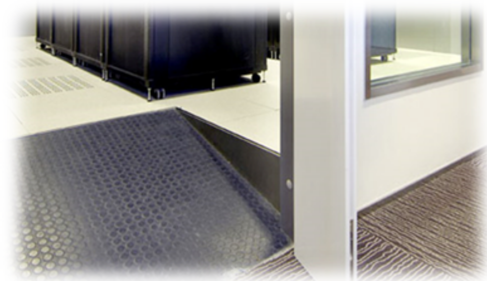
Acceptable Ramp Gradient	Unsuitable Ramp Gradient

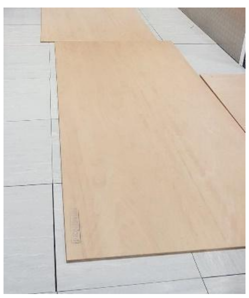	
Spreader plates can be used to increase rolling load capacity	Example of improper floor strengthening

7. IT Technical Space (White Space)(Matrix Ref. 4)

In the OCP Ready requirements matrix, a number of structural attributes for a data center's IT technical 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 uniform load (matrix ref. 4.02), to support a 500kg rack would be ≥ 732 kg/m2 (150 lb/ft2)(7.17 kn/m2). The 'optimum' parameter for the access or solid floor uniform load (matrix ref. 4.02) would be ≥1221kg/m2 (250lb/ft2) (11.97kN/m2).

The 'acceptable' and 'optimum' parameters for the rolling, uniform and concentrated loads have been derived from the performance specifications listed in ANSI/BICSI 002-2019 and The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992.

Fig 7-1

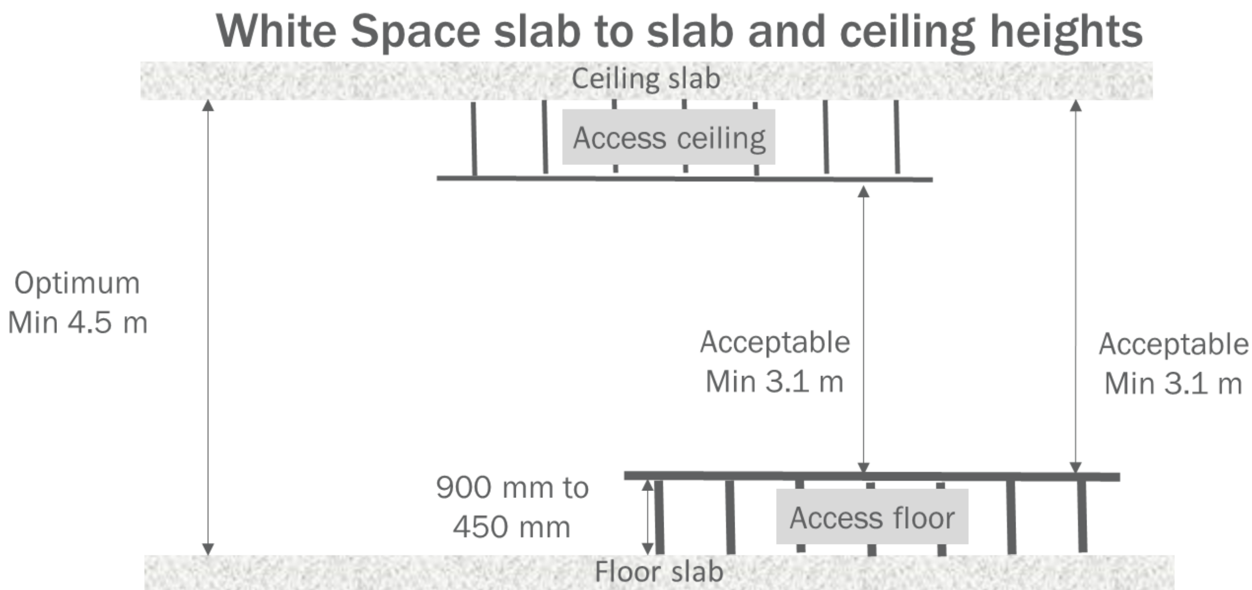

7.1 Pod Architecture

For colocation data centers to meet the demand to scale compute resources more quickly and efficiently in larger increments, pods are being deployed to fulfill this need. A pod is defined as Open Racks installed either in a row or (more typically) a pair of rows, that share some common infrastructure elements such as a RPP, air containment system and cabling pathways.

7.1.1 Pod Framing System Design

A pod frame is a floor standing structure with door frames, doors and panels to enclose a hot or cold aisle zone which contains IT equipment warm exhaust air (HACS) or cooling unit supply air (CACS). Mounting locations are provided for a series of cantilever arms and are able to support a range of infrastructure components, such as cabling pathways for power and network and also for the distribution of electrical and network connectivity via bus bars, tap-offs and patch panels.

Examples of Typical Pod Framing Layouts

Fig 7-2

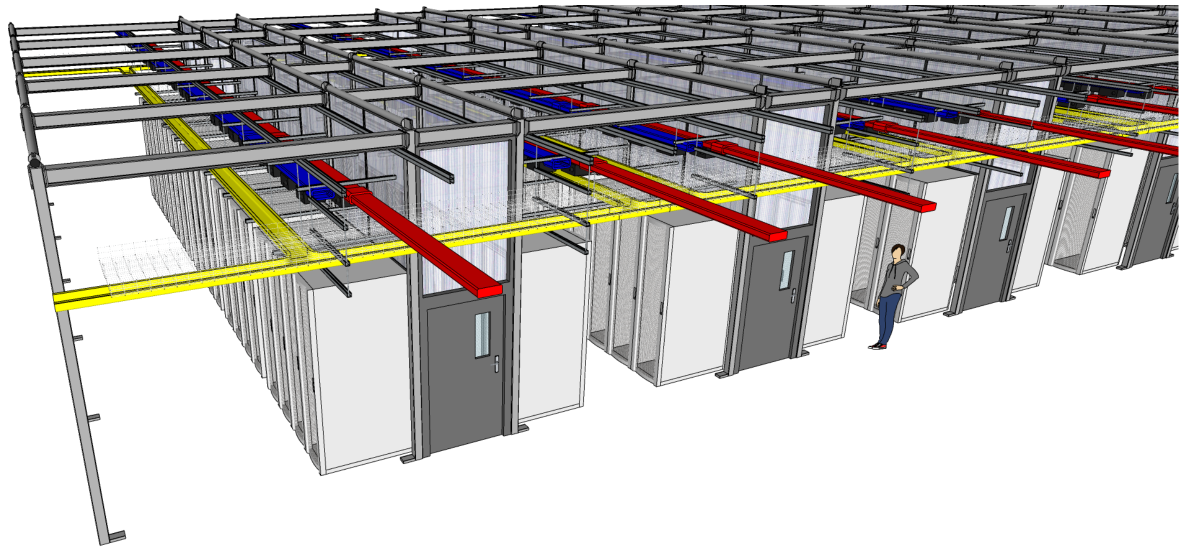

Fig 7-3

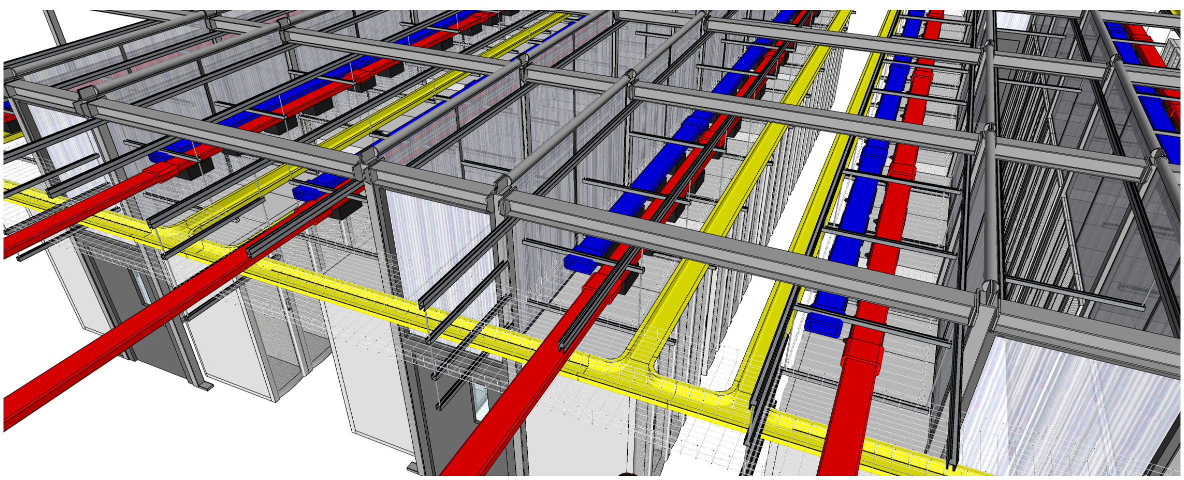

Courtesy of Sean Halpin - Dataplex Group

Fig 7-4

Courtesy of Stephan Harren - Maincubes

8. Electrical Systems (Matrix Ref. 5)

The IT gear within an Open Rack is powered by one or two rack mounted power shelves or PSUs, containing AC to DC rectifiers, which distribute 12V or 48V via busbars in the back of the rack to the equipment.

The Open Rack can also contain Lithium Ion batteries on a shelf (BBU) that would act as the battery backup and therefore providing a CapEx and energy efficiency benefit for a colo to not have to provide a centralised upstream UPS.

For a colocation facility to be able to accommodate an Open Rack that has an IT load of 6.6kW, a ‘must-have’ requirement and ‘acceptable’ level would be to provide a rack supply, fed by a central upstream UPS with a maximum circuit capacity (matrix ref. 5.02) of 3-phase 16 amp, with a power receptacle / WIP type (matrix ref. 5.06) compatible with either IEC 60309 516R6W, NEMA L21-30C or NEMA L15-30C or an optimum level of maximum circuit capacity (matrix ref. 5.02) of 3-phase 32 amp.

The ‘Upstream UPS options’ attribute (matrix ref. 5.07), which has been categorised as ‘optimum’ within the matrix, 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.

Also it is important for an OCP user to understand the generator load acceptance 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 generator 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 ‘acceptable’ 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 Open Rack v2.0 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.

Fig. 8.1

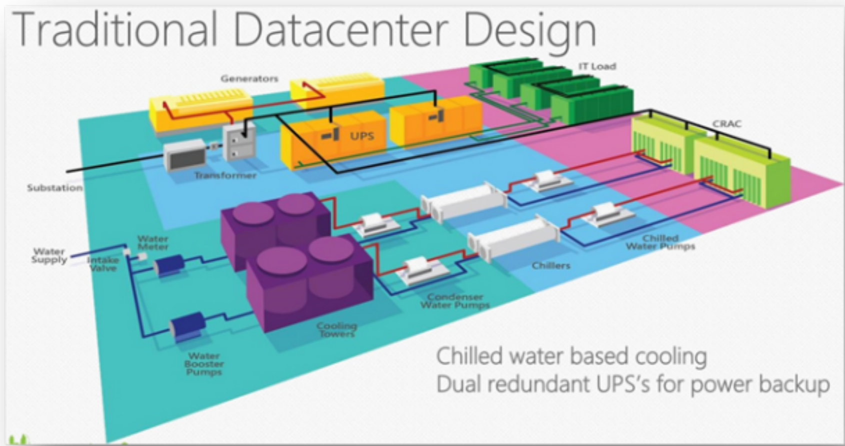

Acceptable Design

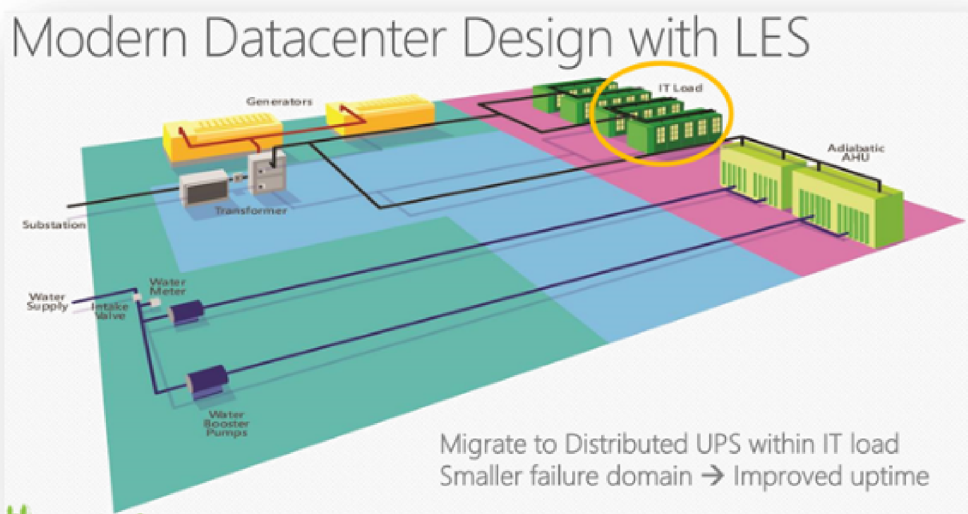

Optimum Design

8.1 Number of independent circuits to the rack - (matrix ref. 5.01)

8.1.1 Open Rack - 1N

Fig. 8-1 below illustrates an OCP-specific 1N (aka BICSI Class F1) power architecture which is the 'acceptable' parameter for matrix ref. 5.01. This design aligns with the simplicity and cost-reduction aims of OCP, and has the following attributes:

- Single path to the IT load
- Minimal breaker count
- No centralized UPS, and instead 1N BBU within the rack
- OCP servers with 1N PSU within the rack

Fig. 8-2

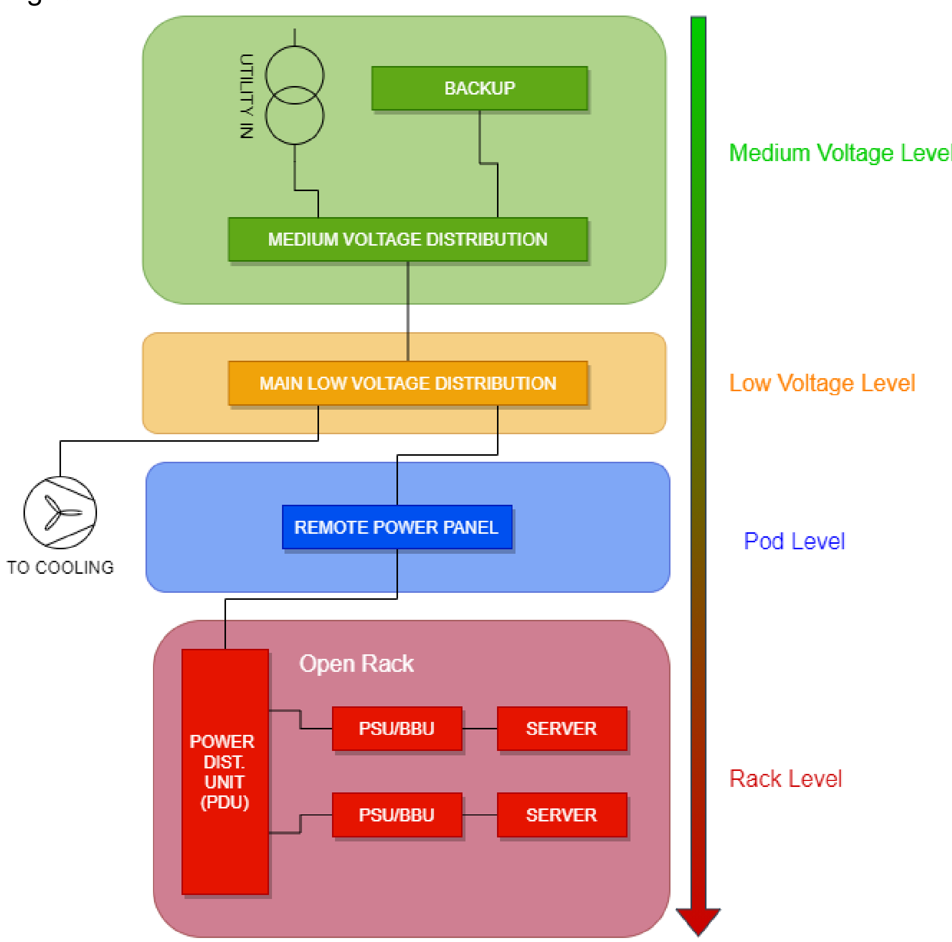

8.1.2 Open Rack - 2N

Fig. 8-2 below is a 2N power architecture for Open Racks only deployed in data center, this design has the following attributes:

- Dual utility feeds with shared backup e.g. generator set
- Dual paths to the IT load
- No centralized UPS
- BBU within the Open Rack (one power path only)
- OCP loads with 2N PSU

Fig. 8-3

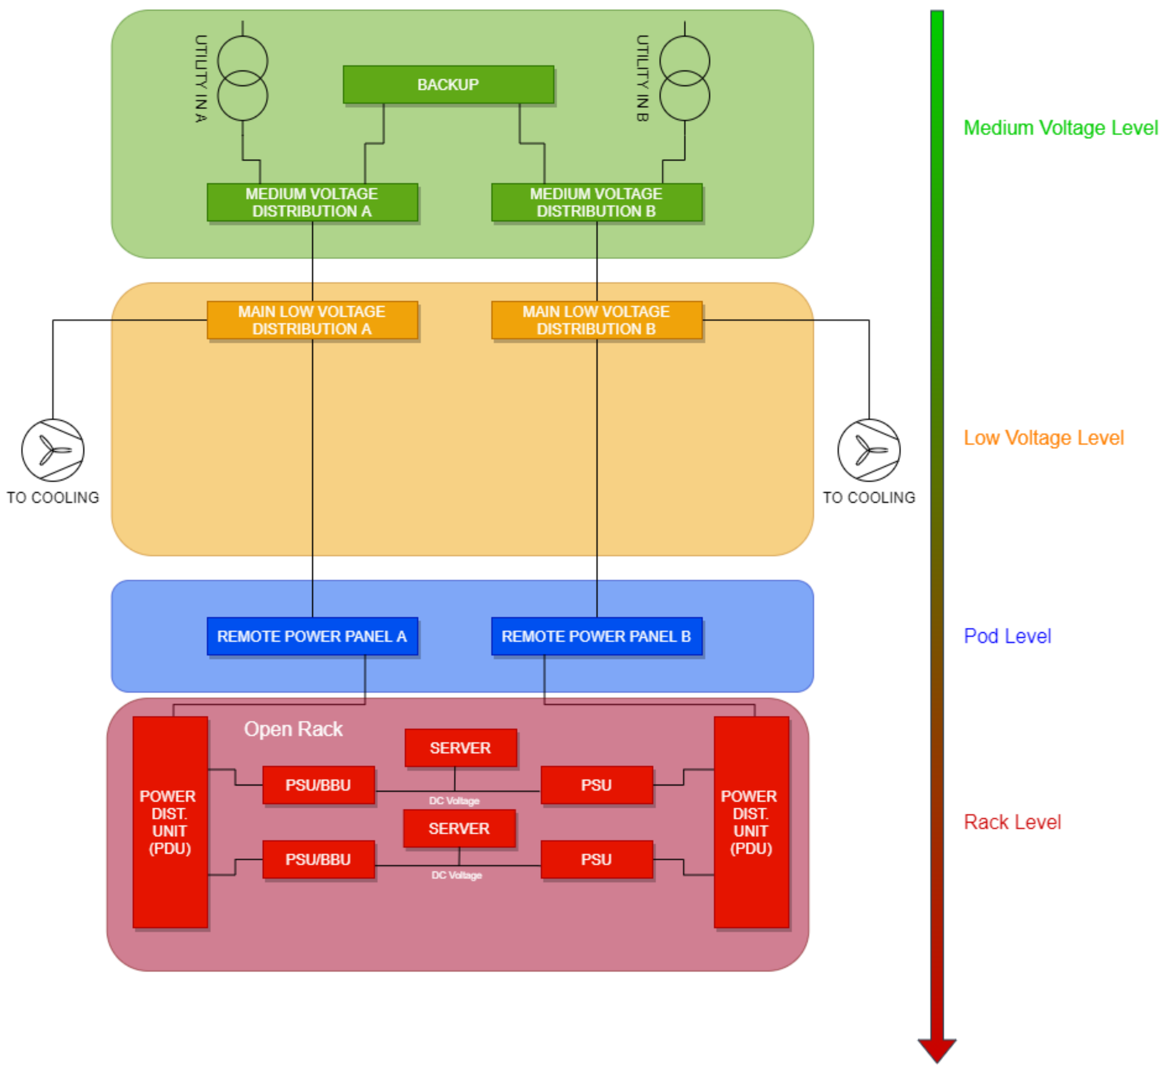

8.1.3 Open Rack & Traditional Rack - 2N

For mixed load environments consisting of traditional and OCP IT gear a 2N power architecture could be adopted. This design has the following attributes:

- Dual utility feeds and backup e.g. generator sets
- A UPS upstream to support both traditional and OCP loads
- OCP loads with 2N PSU and no BBU within the rack
- Traditional loads with 2N PSU

Fig. 8-4

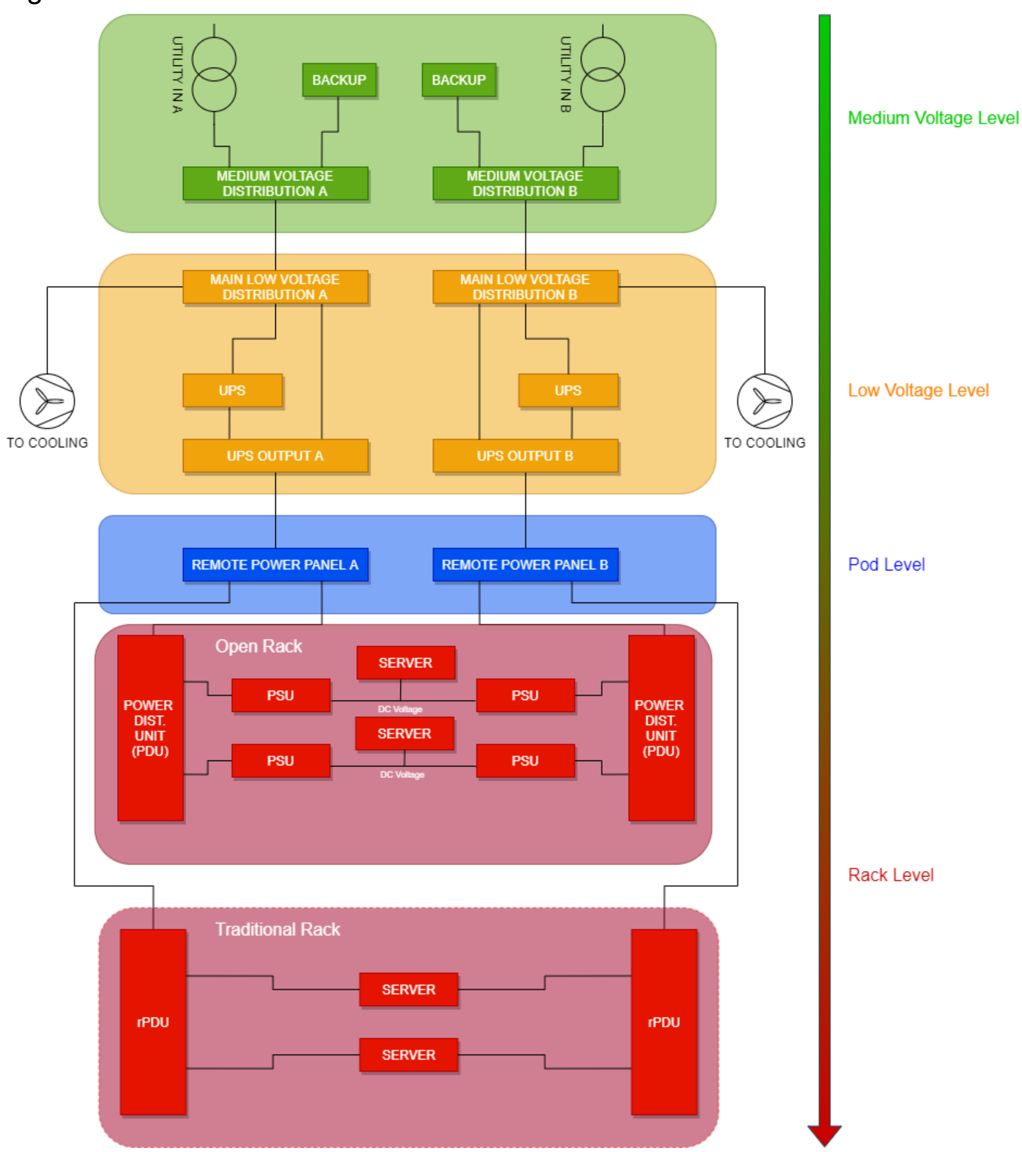

Fig. 8-5 below is an alternative 2N power architecture, this design has the following attributes:

- Dual utility feeds with shared backup e.g. generator set
- A UPS upstream to support both traditional and OCP loads (one path only)
- OCP loads with 2N PSU and no BBU within the rack
- Traditional loads with 2N PSU

Fig. 8-5

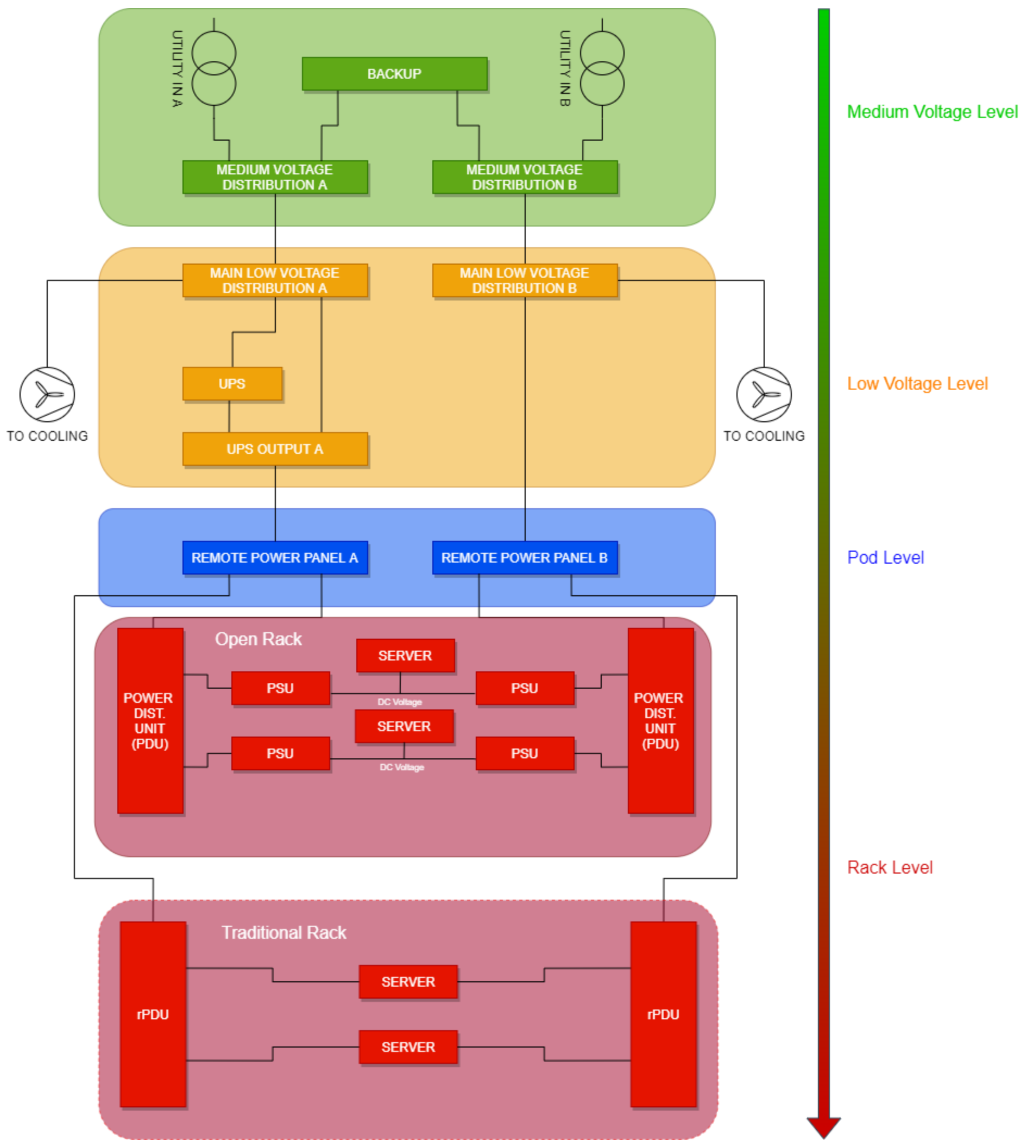

Ref: [Schneider White Paper - WP 228 - Analysis of Data Center Architectures Supporting Compute Project (OCP)]

9. Cooling System (Matrix Ref. 6)

9.1 Introduction

This section of the guidelines deals with the configuration of the deployed racks and the ability of the Data Center Facility (DCF) to deal adequately with the heat load presented by the IT equipment. As in other sections of the guidelines, each subsection relates to the OCP Ready™ Requirements matrix; in particular the attributes that are noted as being important for OCP users of co-location facilities to review and consider during the deployment of their IT installation.

This section does not deal with the design of the cooling system external to the room. Each cooling system will have various merits and features, such as efficiency and resilience. The aim of this guide is not to prescribe the type of heat rejection system, rather it deals with the performance of the cooled air delivery to ensure the user can make an assessment of the facility under review.

Cooling an Open Rack is very similar to cooling a rack of traditional IT gear. Generally, it has front to back airflow and the density can vary from low to high (~4kW to greater than 20kW and potentially even 36kW or beyond in some applications) 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 equipment falls. In the OCP Ready™ Requirements Matrix it highlights for instance in attribute (matrix ref. 6.01), the need for ‘front to back’ airflow. This may well depend on the arrangement of the cooling equipment provided by the data center infrastructure. However, the co-location operator must be able to demonstrate that the facility, when housing an Open Rack or other equipment, is fit for purpose and allows for a managed air flow from a ‘cold aisle’ entering through the front of the server or other IT equipment and exhausting to the rear and in to the ‘hot aisle’ which should be segregated.

To ensure reliable cooling, supply and exhaust air ***<u>MUST</u>*** be segregated to avoid mixing and derating the facility cooling power and reducing the efficiency of the facility. The OCP Ready™ Requirements (matrix ref. 6.02) deals with this must-have.

As a 'Best Practice', air segregation or containment is recommended for any density of deployed IT equipment in order to improve energy efficiency. Perhaps even more important is ensuring the reliable performance of the server or other IT equipment, by ensuring the supply air (SA) is in the warrantable range for the stable operation of the electronics contained within. Requirements (matrix ref. 6.07) covers this issue.

Hot aisle containment (HAC) is considered to be the optimum design, since it meets the goals of the Open Rack design (front and top access during operation) and the 'cool' supply air being supplied in the room void also lends itself to the optimum design of hard floor (solid concrete or similar) facilities instead of access floor/ supply air (SA) void designs.

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 of the rack, so if the rear of the racks are contained in a hot aisle then maintenance personnel will never (or rarely) need to enter that space, which is normally very uncomfortable to work in. Therefore, it is considered that the 'Optimum' arrangement is to have a hot aisle containment system; however, some form of segregation is required to meet the acceptable requirement.

The 'must-have' attributes for a data center facility in this section of the guidelines include:

- IT equipment provided by the user installed in racks must have common direction, front to back airflow – matrix ref. 6.01
- A facility design that supports either hot aisle or cold aisle containment - matrix ref. 6.02
- The ability to deal with the anticipated distributed load densities from 8kW to 12kW and above. This should not restrict the ability to accommodate higher density racks in specific locations. matrix ref. 6.03 refers.
- The IT Technical Space geometry is increasingly important, particularly when considering higher densities. The geometry will impact on the cooling air flow velocity (and pressures) required to meet the cooling power demand of the racks. Therefore, the spacing of racks and relationship to the cooling air flow supply outlet terminals are extremely important. The matrix ref. 6.04, 6.05 and 6.06 all refer to the spacing and

geometry, designed to ensure reliable operation

- IT Technical Space (White Space) environmental conditions that meet current ASHRAE TC 9.9 conditions, relative to supply air (server inlet temperature and humidity). This will be scored as acceptable for the 'Recommended' range or for the optimum efficiency the 'Allowable' limits. The requirements matrix reference for this is 6.07.
- Guideline matrix attribute 6.08 relates to air quality. This is also a very important item to review and consider, since it relates to the potential gaseous contaminants and dust particles that might be found in the data center air flows. This can be affected by the geographical location of the DCF. Its importance is that manufacturers' warranties often specify acceptable limits.
- Temperature rise across the IT equipment is a critical factor to consider and match to the DCF operating conditions and design limits. This matrix ref. 6.09, is described as 'Delta T' for the IT equipment & Open Racks. The sections below will deal with the effect of Delta T and air flow relative to the effective cooling capacity of the DCF.
- Cabinet or rack blanking of open OU space is another simple, but very important attribute to consider. Whilst this is generally a user installed item within the equipment deployment, the co-location operator should be clear in how blanking should be applied to ensure the performance of the data center cooling system meets the design specification and achieves the published efficiency data. This is covered in the requirements matrix ref. 6.10.

9.2 Rack Air Flow Direction – (matrix ref. 6.01)

The Open Rack design is based on a standard air flow and all servers and IT equipment should follow the principle of 'front to back' airflow. This will ensure that the efficiency of the data center is optimal. The aim of this is to ensure the supply air (SA) which should be at a predetermined temperature within the ASHRAE ranges, enters the server or IT equipment via the front face, to ensure that the fans contained within it draw the cooling medium across the critical components 'picking up' heat as it passes through and then is exhausted to the rear.

This process can be hindered if the rack or cabinet contains miscellaneous equipment that has a conflicting air flow direction (such as ‘back to front’ or ‘side to side’), which can then take air from the hot aisle (HAC system or open aisle) and bring it to the front of the rack. The ability of the equipment to do this will be dependent upon the relative pressures in the room, however the pressure differentials across IT equipment are generally low and server fans can easily overcome that pressure.

The result is that warm or hot air will be transferred into the supply air aisle (or containment) and fed back into ‘normal’ front to back airflow devices. This can cause overheating in the worst case, but often, it will reduce the performance of the device.

In addition to the use of equipment with opposing or different air flow directions, it is possible sometimes to install ‘front to back’ airflow devices the wrong way around! The air flow should be determined and verified during the equipment installation.

9.3 Air Containment Methods – (matrix ref. 6.02)

Fig. 9-1 shows alternative arrangements for Open Rack deployments in Pods. Other ‘Proprietary’ IT pods are available, however, the colocation solution provider must be able to accommodate either of the options below, to meet the requirements for an OCP Ready™ data center facility.

The style depicted in a) shows racks founded on a solid floor, with a room air environment being the supply air (SA) path. End doors to the contained hot aisle (HAC) suite of racks are essential to maintain air segregation, particularly at the end rack positions. A ‘chimney’ of insulating and preferably translucent material should be arranged to duct the hot exhaust air to the DCF dedicated hot return air plenum. This should be concealed above the IT Technical Space remaining servicing zone, containing lighting, IT Containment and IT equipment power distribution.

The DCF should ensure that the users HAC design integrates fully with any smoke/fire detection systems provided, fire suppression systems (dry pipe mist or gas) and any other critical systems required for safe and reliable operation of the IT Technical Space.

As an alternative arrangement, but providing optimal performance, individual Open Racks and small groups of Open Racks can utilise a discrete individual chimneys, ducting air directly back to the return air path or plenum. This style of containment usually has a sealed back of rack ‘door’ with a duct or chimney running vertically upwards, reducing the possibility of air leakage and air recirculation.

Fig 9-1 – Air segregation systems

	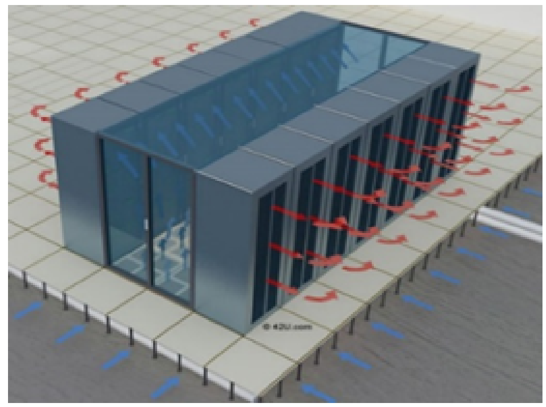
Optimum Design – style a)	Acceptable Design – Style b)

The style depicted in b) shows racks founded on an access floor, acting as a supply air (SA) plenum.

As per the architectural attributes of the OCP Ready^{TM} requirements, the solid floor is the optimum arrangement for OCP racks and standard IT equipment racks, since the raised floor can provide a limitation to the air flow, depending on the outlet grilles (fixed or with directional or flow rate dampers) or congestion under the floor (pipes, power connectors or cabling etc.). As part of the due diligence on a particular facility, where an access floor exists, it is essential to verify that unobstructed air flow can be achieved.

The volume of air available for individual racks can be severely affected as described above; this in turn will potentially limit the power density distribution of the Open Racks in a given IT Technical Space. The contained ‘cold’ or supply air void (Cold Aisle Containment – CAC) may help with air distribution to the racks in any given IT pod framing design, since it forms a constrained 'pool' of cooling air, but it still requires attention to detail to ensure the growing power demands of racks can be met by the available air flow.

End doors to the 'CAC' suite of racks still remain essential to maintain air segregation in order to drive up operational efficiency. However, they generally have less impact on reliability of operation of the racks in the suite of racks due to the volume of cool air available, which then leaks out and bypasses the IT equipment.

9.4 Maximum Rack Density – (matrix ref. 6.03)

Within the requirements matrix, maximum rack density is currently classed in two bands: Acceptable – equal to or greater than 8kW or Optimum – equal to or greater than 12kW.

The available power density within a data center is linked to a number of other interdependent parameters, such as 'Delta T' and air flow rate. In turn, air flow rates will often be impacted by the physical constraints of the data center and the IT Technical Space.

The aim of this section of the guideline is to prompt a review of the power density that the DCF claims to be able to support, as this drives a direct bottom line advantage when procuring the space and power.

It is prudent to review any computational fluid dynamic (CFD) studies that may have been undertaken as part of the design. Typically, these are carried out based on a uniformly distributed IT load. They should clearly demonstrate the ability of the cooling arrangement to support either of the density ranges stated as a minimum requirement.

Where the user might have a non-uniform power density, the results of the CFD study must be interpreted or re-run to get assurance that higher specific rack densities can be supported within the overall kW load allowance for the zone of the facility. For instance, whilst a row of 12 racks, rated at 12kW (144kW total) might be supported well, the facility may not support 3 racks of 36kW (sub-total of 108kW) and a balance of 9 racks at 4kW (sub-total – 36kW).

9.5 Cold and Hot Aisle Widths – (matrix ref. 6.04, 6.05 & 6.06) (IT Technical Space Geometry)

As discussed in the preceding section, the ability to review a CFD study for a particular data center facility is extremely useful. Whilst it is not a formal requirement for the OCP Ready™ program, it is a great tool to assist in the assessment of the facility and if it is 'fit for purpose'.

A basic CFD output should indicate a uniform velocity to all parts of the IT Technical Space where load is deployed. Depending on the cooling air flow configurations (one side or both sides of a room, top down etc.) velocities may be higher at air flow discharge points; however, this should quickly disperse and create a uniform velocity. The aim of this guide is not to be prescriptive but provide background information in order to allow the colocation facility user make a reasoned choice in relation to deploying an IT platform within a particular DCF.

The spacing of racks will be dependent upon rack power density. Large groupings of kW capacities above a uniform design density may create a demand for high velocity air due to size constrained aisles or corridors. Basic analysis can be undertaken by considering the cold and hot aisle widths and the obstructions in an aisle or SA delivery path.

Planning should be undertaken to develop the standard layout of an IT Technical Space pod (similar in style to that shown in fig. 9-1 a) or b)) and set out in a plan view drawing similar to that shown below in fig 9-2. An assessment of the rack count, pod spacing, and kW allocation can then be undertaken.

Fig. 9-2 – Typical plan view of an IT Technical Space layout

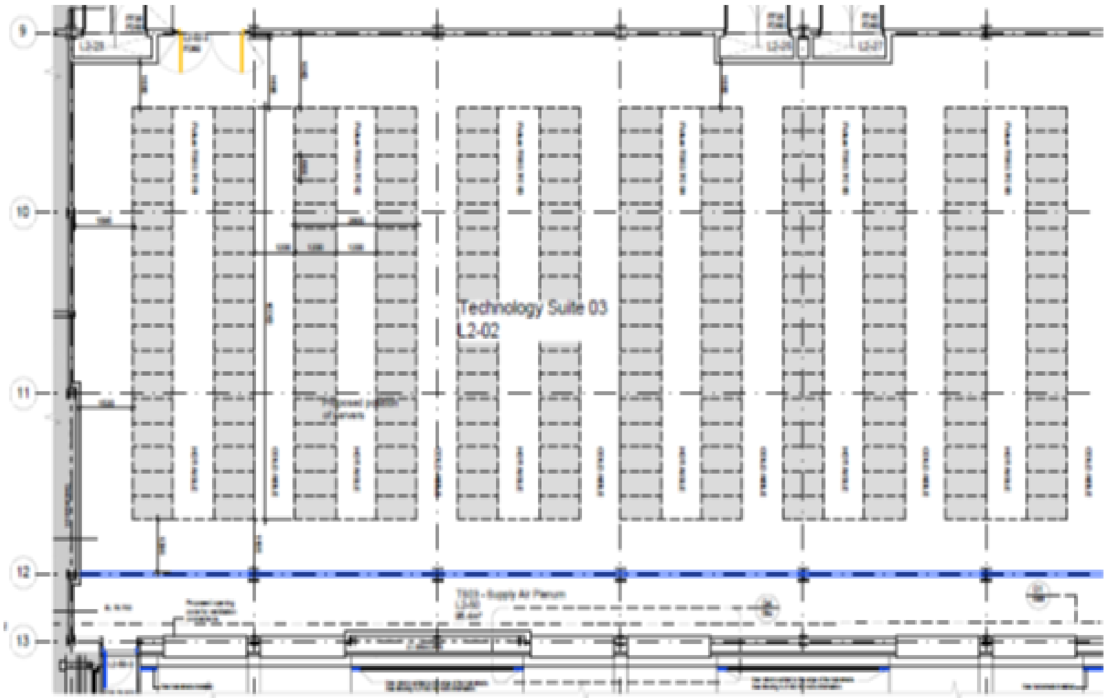

Using the guidelines, the assessment can be made in a basic form by looking at the aisle widths; -

- Minimum Cold Aisle width - (matrix ref. 6.04)
  - Acceptable Parameter — Greater than or equal to 1200mm (48")
  - Optimum Parameter — Greater than or equal to 1500mm (60")
- Minimum free width - Cold Aisle (inside a cage) – (matrix ref. 6.05)
  - Acceptable Parameter — Greater than or equal to 900mm (36")
  - Optimum Parameter — Greater than or equal to 1200mm (48")
- Minimum Hot Aisle width – (matrix ref. 6.06)
  - Acceptable Parameter — Greater than or equal to 900mm (36")
  - Optimum Parameter — Greater than or equal to 1200mm (48")

In reviewing the first parameter (matrix ref. 6.04), this is a measure of how easy the air flow will make its way to the IT equipment, housed in the rack. At the densities considered in the guidelines, typically, 1200mm (48"), or approximately the same width as the depth of a typical rack, is generally acceptable.

Where the IT power density is higher, then a greater width will offer less resistance to the required air flow; however, if the server fans are strong enough to draw the air in, then the result in a narrow cold aisle is that the air velocity increases.

The detrimental effect of this is only really seen when a CFD study is reviewed. The tendency is, particularly where cooling air is introduced on one side of the IT Technical Space only, then at the end closest to the cooling air supply, the air velocity will increase the most. If the velocity becomes too high, it can bypass cabinets since the velocity is higher than the suction of the server fans, which can result in a negative pressure in front of the IT equipment. At its worst case, this could cause hot air from the contained hot aisle to be drawn through leaky rack sealing provisions and then allow the hot air to be drawn back into the server air inlet. Eventually this could result in overheating.

The extract from the CFD study on the following page (velocity and pressure) show this phenomenon.

Fig. 9-3 – Vertical velocity plane in cold aisle – end unit

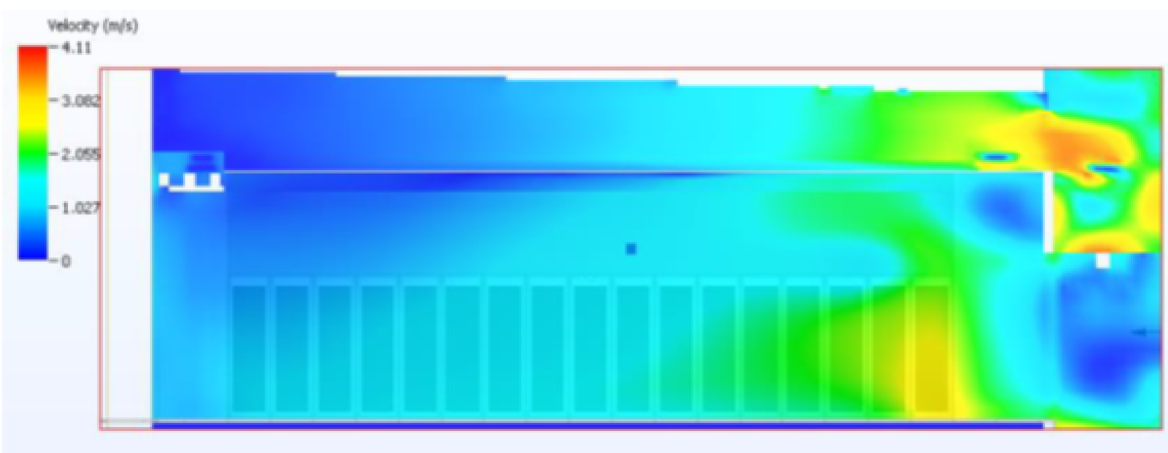

Fig. 9-4 – Vertical pressure plane in cold aisle – end unit

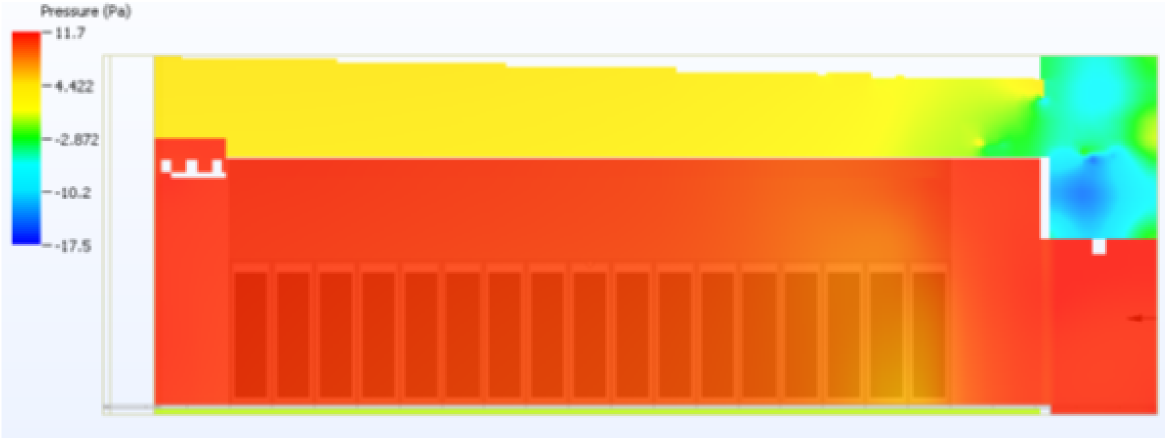

It can be seen that at the end of the cold aisle, closest to the cooling air supply, the velocity is at a velocity of between 2 to 2.5 m/s (green/yellow shading). At this speed the solution performs well, and racks are not starved of air, nor are significant reverse or negative pressures developed, forcing air back to the cold aisle from the hot aisle. However, the pressure plot does indicate a reduced positive pressure of say 4.5 to 6 Pa (yellow/orange shading), compared to the surrounding areas where it is uniformly 11 Pa (red shading) or higher.

If a CFD shows velocities significantly higher than this and certainly when over 3.5 to 4m/s, a more detailed study, particularly in cooling redundancy failure modes should be requested. If velocities are too high, the ability for fire detection equipment to work correctly can also be affected. It would also be likely that the reduction of pressure in that case could start to make the HAC to cold aisle pressure negative, allowing bypass air to leak through.

The parameter of the second attribute (matrix ref. 6.05) relates to a minimum width again for a cold aisle. This is not a contradictory statement, even though it specifies a smaller dimension. It is more a realisation of a physical constraint for IT equipment handling. Typically, the cages used in an IT Technical Space are used to segregate different customers' equipment. A cage is usually open in construction to allow air to pass freely, in order to maintain cooling effectiveness.

However, in dividing space, to ensure the colocation operator does not lose too much rentable area, there is a desire to minimise circulation space sizing. It is important that when installing an OR, whatever space is left between the cage and the OR, it is large enough to enable the IT gear e.g. server etc. to be fully withdrawn from the Open Rack, to enable replacement, for example. The standard length of the server can be facilitated within the ‘acceptable’ width dimension; however, the optimal dimension allows more space to manoeuvre the server and use a server lifting trolley. This is especially important for heavier and deeper (larger OU) size units.

The last attribute in this section is the hot aisle width (matrix ref. 6.06). Typically, this can be narrower than the cold aisle, but should always be in proportion, since it is being used to transport the same volume of air as the cold aisle, but higher velocities can be tolerated. This is because the cooling equipment air movement fans are drawing the warm air away and it is not as important to the exhaust air that a uniform temperature reaches its destination, as long as the heat is dissipated. Conversely, if the cold aisle velocity is too high as demonstrated above, this can cause overheating issues for the IT equipment.

Whilst a rough guide can be calculated from the anticipated air flow volume relative to the cross-sectional areas of various aisles, it only gives a uniform approximation. A CFD study is the most accurate way of confirming that a given DCF will meet the air flow/velocity/power density and supply air temperature distribution required for an IT equipment deployment. In all cases, providing wider aisles for both cold and hot air will undoubtedly help with the cooling process; however, space costs money and to achieve the power densities outlined in the guidelines, the spatial parameter set out should be adequate,

As power densities increase, or where there is an uneven power distribution, then additional space may be required.

9.6 Inlet Air Conditions - (matrix ref. 6.07)

The majority of IT equipment is manufactured in line with the environmental conditions specified within the output of the ASHRAE Technical Committee 9.9. The current version was issued in 2015 and is set out below. Updated requirements should be reviewed as they become available.

There are two principal operating conditions that are generally available within OCP Ready™ colocation data center facilities and they are: -

- ‘Recommended’ – the Class A1 envelope, providing a supply air range of 18 deg C (64.4 ⁰F) to 27 deg C (80.6 ⁰F) and humidity in the range of approximately 5.5 deg C (42 ⁰F) Dew Point to 60% RH and 15 deg C (59 ⁰F) Dew Point.

- ‘Allowable’ – the Class A1 envelope, providing a supply air range of 15 deg C (59 ⁰F) to 32 deg C (89.6 ⁰F) and humidity in the range of 20% RH to 80% RH.

These environmental ranges are set by ASHRAE in conjunction with inputs from all major server manufacturers. Specific Open Compute Servers may have different ranges but will be aligned to these limits. The chart below shows the temperature operating zones.

Fig 9-5 – Typical ASHRAE Recommended and Allowable Zones on the Psychrometric Chart

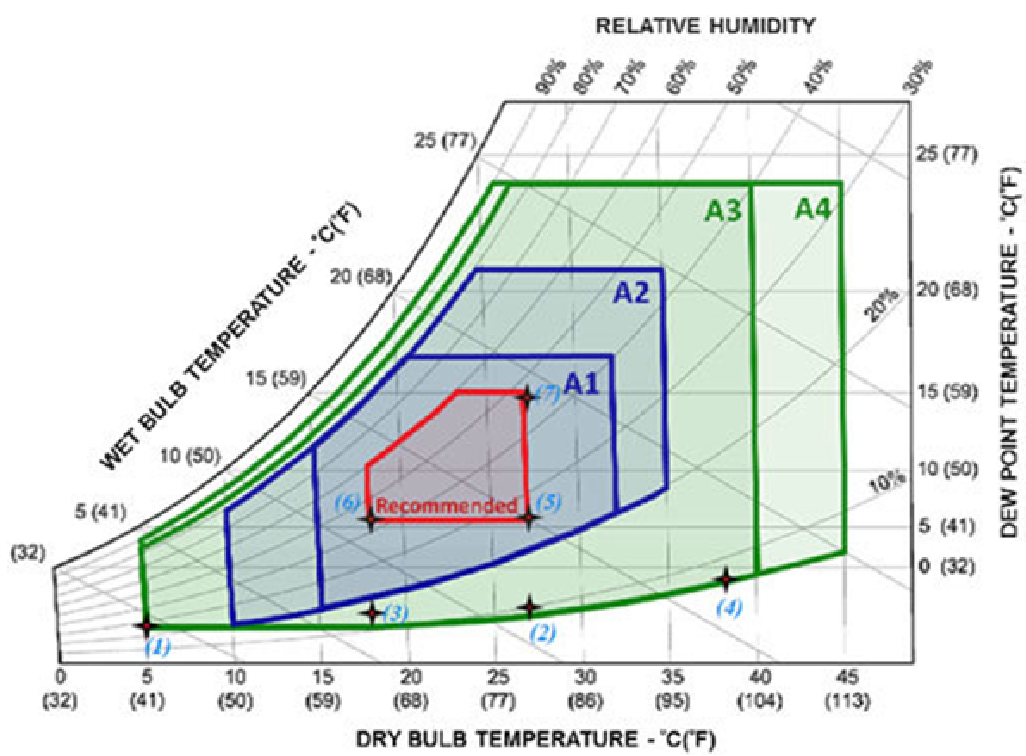

-

In the requirements matrix, the optimum operating range for attribute matrix ref. 6.07 is classed as the ‘Allowable’ envelope. The reason for this is that at peak outdoor ambient temperatures or ‘Summer Conditions’ (certainly in the Northern Hemisphere), the cooling system will be more efficient overall by not having to overcool the air supply and that will drive a lower PUE, since the cooling system is not trying to reject more energy (heat) whilst the external ambient is high. In most colocation facilities, energy is charged using PUE as a factor, so a reduced PUE will also translate into lower operating costs for the end user.

The key attribute here is however, that OCP server and IT equipment is generally designed to operate reliably within the allowable limits (up to 32 deg C or 89.6 ^{0}F) and often beyond up to 40 deg C (104 ^{0}F). It is sensible to verify which IT equipment platform is to be deployed, but providing its inlet air requirements sit within the ‘Allowable’ range, then stable and reliable operations should be achieved.

Where the IT equipment is not rated to operate up above say 27 deg C (80.6 ^{0}F), it may be necessary to accept a tighter temperature operating band. This is called the ‘Recommended’ range and is generally an envelope of operating conditions that applies to a wider range of equipment including some older ‘legacy’ equipment. As noted above, specifying this range is likely to result in increased operational costs at the data center facility, so where possible, OCP hardware that can operate beyond the ’Recommended’ range should always be the first choice.

Temperature operating or environmental conditions often form part of the server manufacturers warranty conditions. If equipment becomes damaged through operating beyond the set limits may invalidate the warranty.

Additional warranty conditions may be imposed by the server manufacturer and the operator should check against other requirements, such as gaseous contaminants in the IT Technical Space etc. This topic is covered in the requirements matrix in the next section, matrix ref. 6.08.

9.7 Air Quality – (matrix ref. 6.08)

The previous section of the guidelines dealt with the operating temperatures applied to the operating conditions for OCP hardware as well as other IT Gear. This section deals with the quality of the air and the effect that it can have on the reliable operation of the servers.

Firstly, this section of the guidelines is concerned with the physical aspects of air quality and relates specifically to airborne particles. The source of the particles can be from personnel at the data center entering the IT Technical Space and bringing fibres and dust in from external areas, on items such as clothing and IT equipment that has been unpacked or stored in a non-clean environment.

To try and prevent dust particles entering from other sources, such as airborne dust in the surrounding circulation spaces leading to the IT Technical Space, typically a data center facility will have what is termed as a 'Fresh air make-up AHU' or 'Pressurisation AHU'. This provides fresh air ventilation to the space and the aim of that plant is to pressurise the IT Technical Space with fresh air, such that the internal pressure stays at say around 20 Pascals, relative to the surrounding rooms and corridors. By doing this, the tendency is to leak air from the room, preventing dust being drawn in through small openings and cracks as well as keeping out the surrounding circulation space air when doors are opened in the process of accessing the room.

To address both sources of particulate contamination, it is important that the data center operator can demonstrate that: -

- It has adequate air filtering to EN 779 G4 and F7. These are 'General' and 'Fine' filter classes and could be included in:
    - Any internal recirculating cooling system to scrub the air and remove particles brought in by personnel etc.
    - External air AHU used to pressurise the space, to reduce the risk of outside air particles reaching the IT Technical Space.

The data center facility specification and drawings should confirm what has been included, to enable the user to verify this part of the OCP Ready™ requirements.

In addition to the basic particle and dust contamination, a lot of work has now been carried out by a range of server manufacturers in relation to the effect of airborne gaseous contaminants. This has become a critical attribute in the assessment to observe, since it again feeds into the underwriting of the IT equipment warranty. The standard that applies to the monitoring of gaseous contaminants in the air is the ANSI/ISA 71.04 - 2013 standard, which refers to a series of contamination severity levels. These carry the nomenclature G1, G2, G3 and GX. They cover the level or risk of the gaseous contamination as 'Mild', 'Moderate', 'Harsh' and 'Severe' respectively. Each 'G Level' refers to a particular level of corrosion in either copper or silver metals. The rate of corrosion is measured in Angstroms, fig. 9-6 below shows the ranges.

Fig 9-6 – Gaseous Contaminant Levels for Copper & Silver

Severity level	G1 - Mild	G2 - Moderate	G3 - Harsh	GX - Severe
Copper reactivity level (in angstroms)*	< 300	< 1000	< 2000	≥2000
Silver reactivity level (in angstroms)*	< 200	< 1000	< 2000	≥2000

*Normalized to a 30-day exposure. 1 angstrom = one hundred-millionth of a centimeter, or 10^{-10} meter.

The importance of this monitoring has increased over the years, with the abolition of lead solder, which was a more stable and less reactive metal. Lead was used in solder on circuit boards and was removed due to the harmful health effects it can have.

Lead solder was replaced by silver-based solder and this is used in the mounting of electronic components in IT equipment, connecting the devices to the copper circuit 'tracks', that interconnect the processing systems. It is therefore important to monitor what the atmosphere contains, to ensure that the data center facility does not put the soldered joints and copper tracks at risk, due to corrosion, thus potentially causing the IT equipment to fail prematurely.

Most server and IT equipment manufacturers offer a warranty on the equipment, providing it has been operating within a G1 environment. Therefore, it is important to validate that a data center can monitor the levels of gaseous contaminants adequately by using reliable equipment. In addition, as part of procurement due diligence, it is recommended that the colocation solution provider can demonstrate how a contaminant level above G1 could be detected and then subsequently dealt with, typically by bio-chemical filters etc.

Therefore, in terms of the OCP Ready™ requirements, the optimum arrangement is to have the appropriate test and monitoring equipment (matrix ref. 6.08)

9.8 Temperature Rise (matrix ref. 6.09) - (Relationship to Air Flow, Power Density and Velocity)

The data center facility will typically be designed around an industry standard for Delta T across a rack. In the data centre industry, this is commonly set at 12 deg. C (53.6 ^{0}F). It is the key critical parameter that customers must be aware of and review in relation to their deployments, since there is an inextricable physics based link between air flow (velocity and volume) and the heat removal from the server equipment.

It is important to ensure that the IT Technical Space is suitable for the IT equipment proposed for the deployment, so as a ‘must have’ the DCF should state the Delta T, in order to make an assessment of the facility relative to the proposed IT equipment deployment. However, if there is a mismatch, there is potential for a shortfall of cooling air flow when the upper limit of the electrical power capacity (kW) is reached. As a ‘nice to have’, the Delta T should be a value 12 deg. C (53.6 ^{0}F) for an optimal score. As shown in the guideline’s matrix, it is possible to have a design based around a Delta T of 8 deg. C (46.4 ^{0}F) for an acceptable result. However, this is likely to be less efficient and will certainly require higher air flows to achieve the same amount of heat removal from the IT equipment.

Depending on the equipment, higher efficiency and hence reduced operating costs may be achieved with a higher Delta T, since the volume of air required is reduced, so DCF cooling system fans are able to work less hard.

If IT equipment with a Delta T of less than that of the DFC design condition, the case may be that the cooling equipment is not able to produce adequate air flow and the result being that the IT equipment overheats.

Older IT equipment can have a smaller Delta T operating point, but even newer equipment can be set to operate at a low Delta T, since it may increase the performance of the system or reduce the potential of damage to the components inside the server, for example. It is important to check with the manufacturer, and check if BIOS settings can be altered to achieve a desired Delta T. The critical attribute to achieve is the highest safe working Delta T, BUT ensure it matches or exceeds that of the data center cooling system.

As the title of this section suggests, another critical issue to plan for, is the deployment of the power density and review air flow velocity, but in conjunction with the Delta T parameter. They are linked and a variance in one, may have a significant effect on the cooling outcomes.

The graphs below show for varying values of Delta T, what air flow is required in order to reject the heat from a rack. As it can be seen, for a given Delta T, there is a linear relationship between air flow on IT load in kW, such that if the air volume is halved, so is the kW of heat. Therefore, it can be deduced that if a server is producing a Delta T of say 7 deg. C (44.6 ^{0}F) only, rather than being matched to a DCF design of 12 deg. C, then the server will require approximately 1.2m^3/s to cool 9kW, instead of the allocated 0.7m^3/s. The longer term effect of that is that if the same pattern is repeated over and over again for each server, the DCF cooling system will not be able to deliver the correct volume of air and eventually, there is a risk that IT equipment will overheat.

Fig. 9-7 – Sensible Cooling Load vs. Air Volume

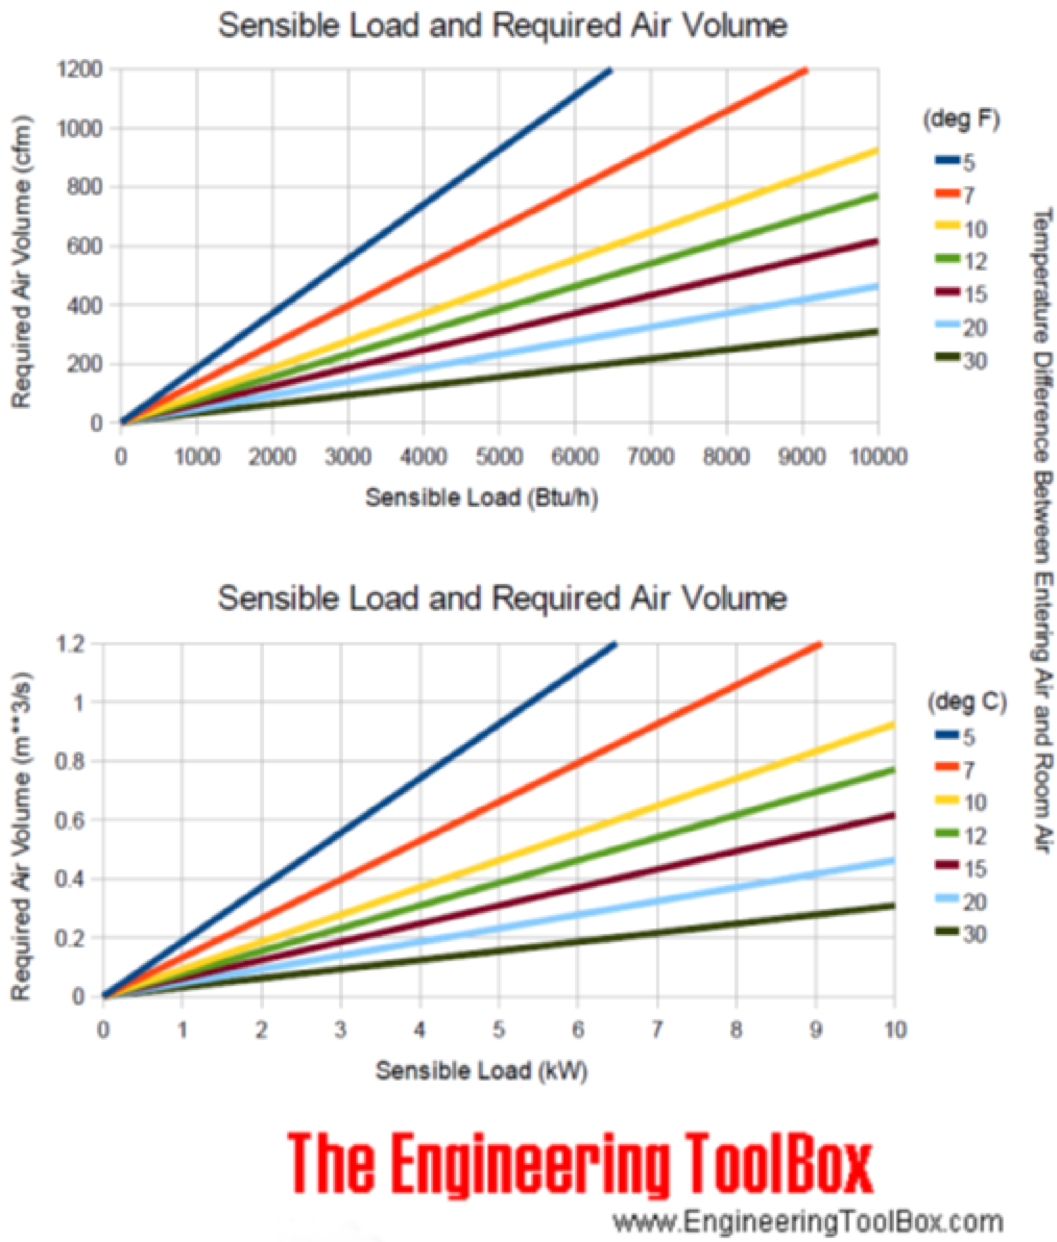

As a ‘nice to have’ condition, the data center facility should be able to demonstrate modelled examples of air flow and velocity vectors (using Computational Fluid Dynamics – CFD) to help customers plan the power density of cabinets. Most facilities will be designed for uniformly distributed IT loads of ‘X‘kW per rack. However, this even distribution of power rarely occurs.

The relationship of velocity, air flow and heat rejection for a given Delta T are discussed earlier in this guideline section, under attributes matrix ref. 6.04 to 6.06.

9.9 Cabinet Banking of Open Space – (matrix ref. 6.10)

This is really the last piece in the puzzle and is a critical area of concern. As the sections above discuss, there are direct links between the key parameters involved in the effective and efficient cooling of IT equipment within a data center.

We have seen how spatially constrained air delivery routes (cold aisles) can cause high velocity air flow, at the air volumes required to meet a specific cooling demand or power density, based on a particular Delta T.

It has also been described how the segregation of the cold and hot aisles is very important to help drive up efficiency and avoid hot exhaust air from one piece of equipment being recirculated through neighbouring IT equipment, therefore creating the potential to cause overheating and thus, unreliable server operation.

In the acceptable category for this feature of the guideline, it is ‘Recommended’ that blanking for every spare space in the rack is fitted. This can be applied to cable ways as well as spare server or equipment mounting spare ‘OU’ slots.

For the IT Technical Space to work in its ‘Optimal’ mode, then effective blanking is deemed as mandatory.

The directly related reasons for this are set out below: -

- Where the air is moving at a high velocity with the cold aisle of the IT Technical Space, it is possible that a negative pressure relative to the hot aisle can be created in the cold aisle. If blanking is poorly fitted or even omitted, then the hot exhaust air in the contained hot aisle will travel through the un-blanked opening and either mix with the cold aisle SA, or directly get drawn into a specific item of IT equipment. This can be detrimental to that equipment and may in time cause it to slow down (in terms of processing speed) or even fail.

· Where the air in the cold aisle is positive in pressure relative to the hot aisle (normal situation, although by a marginal amount only), then the missing blanking plates can create a bypass route for the cooling air (SA). Therefore, rather than performing 'real work' by cooling the intended receiving IT equipment, that air is just blended with the hot air and returned to the cooling system for re-processing. This effectively wastes energy. If the blanking is carried out diligently, then only the right amount of air is delivered to the space and the cooling system has to work less hard.

So, in conclusion, blanking all open spaces should be seen as essential, and in any facility, the optimal condition should be targeted. This will save more money by increasing the efficiency of the cooling system and avoid the potential of overheating IT equipment and reducing equipment life and causing unstable operation.

10. Telecommunication Cabling, Infrastructure, Pathways and Spaces (Matrix Ref. 7)

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, Open 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 'cabling infrastructure pathways' (matrix ref. 7.01) and the routing of network cabling into an Open Rack would be either top or bottom entry and to the front of the rack. The v2.0 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 for the data center is very much specific to the needs of the OCP user's use case. Attributes to be considered by the user 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.

Fig 10-1

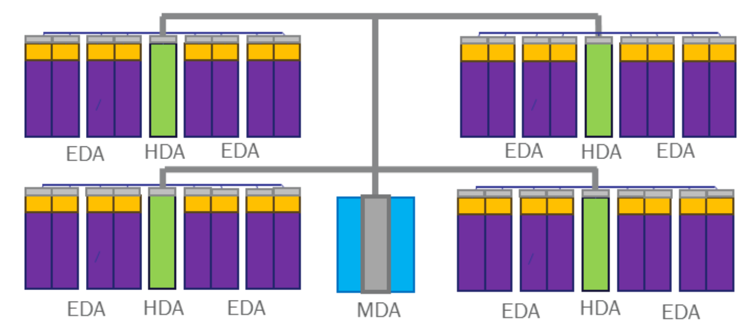

Consideration for Space Design

Fig 10-2

Optimum Pathway Design

11. 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 Technical 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 Open Racks either in a row or (more typically) a pair of rows, that share some common infrastructure elements such as a PDU, air containment system and network.
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 Technical 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.

12. 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 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

13. Appendices

Appendix A. 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.

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.

Appendix B. Sustainability

It is becoming increasingly apparent that to achieve true sustainability for the IT industry, focus should not only be on the energy consumption of a data center but more on greenhouse gas emissions. Colocation data center operators and OCP users also need to consider the wider environmental impact of their IT infrastructure, including aspects such as electronic waste, water usage and the materials used in the construction of their data centers and IT gear.

One of the mitigating actions supported by OCP is the use of a circular economy approach to sustainability, which minimizes waste and makes the most of IT assets. This approach focuses on developing products so they can be reused and repurposed with no residual waste.

B.1 Greenhouse Gas Protocol (GHG Protocol)

Fig. B-1

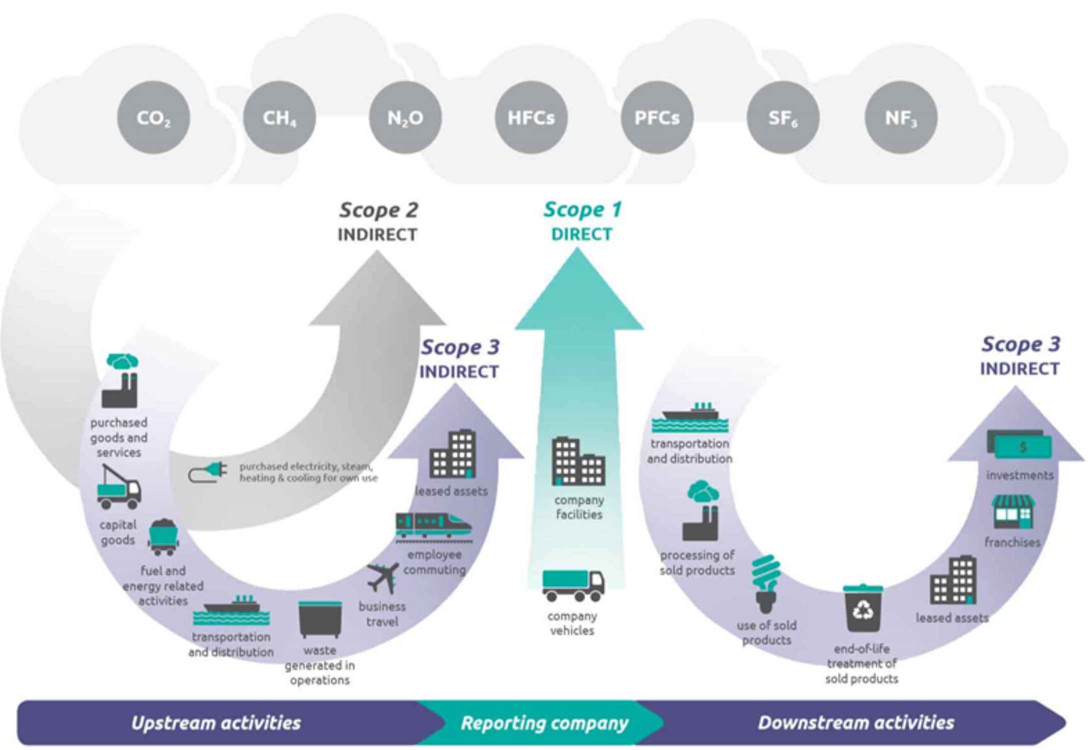

Ref: https://www.opencompute.org/circular-economy

Appendix C. Open Compute Open Rack Types

There are several variations on OCP IT gear, so it is important that the operator of a colo facility works directly with the OCP user to understand exactly what will be deployed and even guide the choices of OCP configurations to best suit the data center and the IT power requirements.

Below is a table showing the two main Open Rack types and their fundamental specifications.

Attribute \ Rack Type	**Open Rack v2.0**	**Open Rack v1.0**
outside width	600mm typical, but can vary	600mm typical, but can vary
depth	Standard: ~1048mm Shallow: ~ 762mm	Standard: ~1067mm
height	unspecified, but facebook uses 2210mm	unspecified, but facebook uses 2210mm
weight load of equipment	variable between 500 - 1400kg	Standard: up to 950 kg
mounting rail spacing	21”	21”
U Spacing	48mm OpenU (OU)	48mm OpenU (OU)
Required access	primarily front only	primarily front only
PSU architecture	3 phase AC rack PSU to 12V or 48 VDC 1 busbar distribution	3 phase AC rack PSU to 12V 3 busbar distribution
Battery Backup	optional in-rack Li-ion	optional in-server Li-ion
Power feed to rack	Typically 3 phase AC 230/ 400 VAC	Typically 3 phase AC 230/ 400 VAC
Airflow	front to back	front to back

Image		
Specification	Open Rack v2.2	Open Rack v1.0 (Facebook)

Appendix 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 the 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 each rack or the whole cable plant.

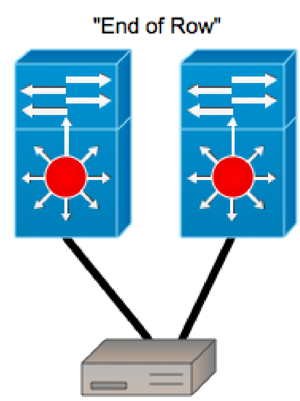

3. Centralized Distribution: This could be the most expensive option but it is the most flexible. This layout ensures that you are ready to receive whatever the data center has to throw at you. Cabling will be installed above the racks attached to the cable tray. The cabling would traverse the cable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

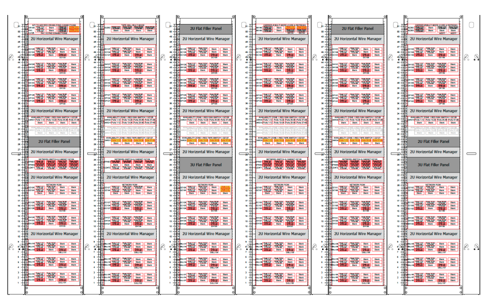

be laid out to make a map of the room. The patch field would also give you the ability to jump from rack to rack, row to row, and room to room. Commonly used with top of rack but will support U network topology as well.

Fig D-1

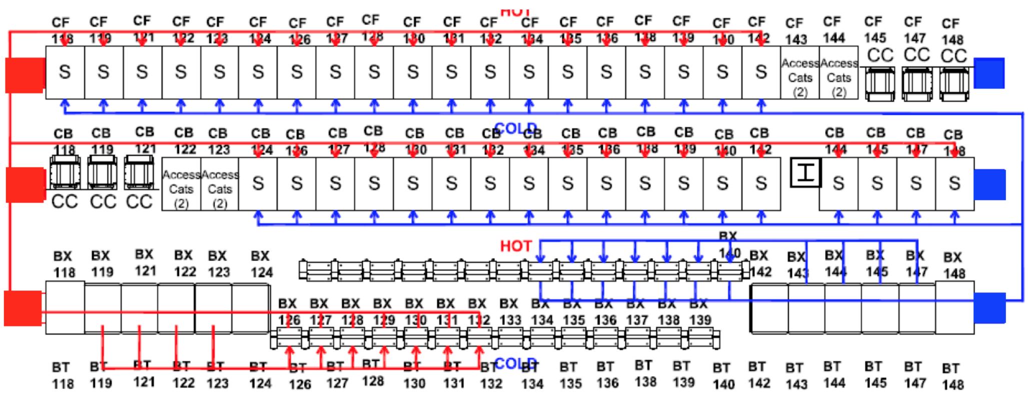

Appendix E. Cable Plant Reference Architecture

E.1 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.

E.2 Leaf and Spine Fabric

Fig E-1

2K Array" Example

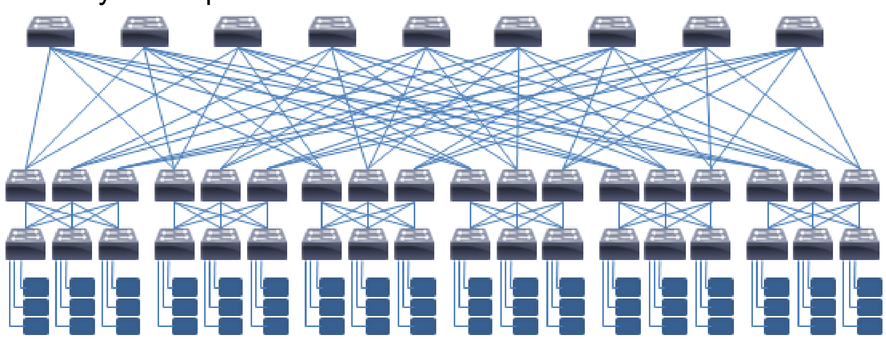

- Many leaf and spine variations are possible. One example is the "K Array" shown here for reference.
- "Fat Tree" characterised by same bandwidth at each layer.
- Folded mesh allows extensive scaling with equal delay connections.
- Non-blocking fabric assuming uniform packet distribution.
- Single switch 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.

E.2.1 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

E.2.2 Leaf/Spine Physical Design 1024 Servers Example

Fig E-2

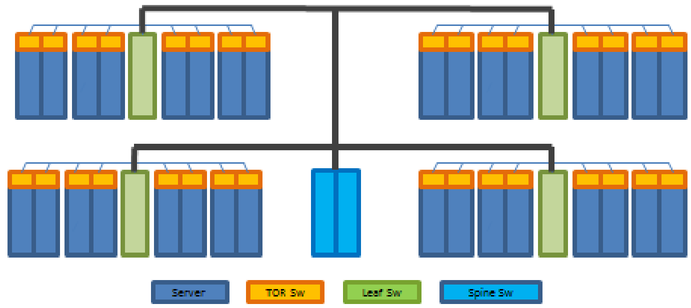

Example

- 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]

E.2.3 Structured Cabling View 1024 Server Leaf/Spine Example

Fig E-3

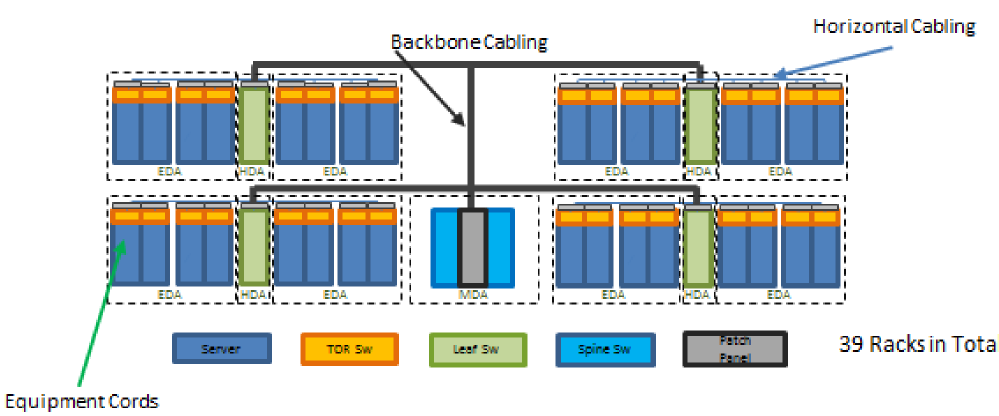

Equipment Cords *10G and 25G* *DAC, Cat6A, MM*	**5120**	**EDA Patch Panels** *Duplex LC, LC/MPO Modular*	**32 x 32 Ports**
Horizontal Cables *AOC, MM, SM*	**32 x 32 Channels**	**HDA Patch Panels** *Duplex LC, LC/MPO Modular*	**4 x 512 Ports**
Backbone Cables *SM, MM (smaller configurations)*	**4 x 256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

Appendix F. OCP Ready Requirements Matrix

Use Information

A. The 'must-have' ‘Acceptable’ requirements are those that a colocation operator would need to implement to enable the trouble free and smooth implementation of an Open Rack (OR) v2.0 that is populated and weighs 500kg and has an IT load of 6.6kw. This is also the minimum requirement for OCP Ready™ certification.

B. The 'must-have' ‘optimum’ requirements are those that a colocation operator would need to enable the implementation of an Open Rack (OR) v2.0 that is fully populated and weighs 1400kg and has an IT load of 12.5kW. If all ‘Optimum’ requirements are met by the colocation operator, it would enable the full benefits of the Open Rack design to be achieved alongside being certified OCP Ready™.

C. There is also additional guidance information in the requirements matrix in the Electrical section, for example, the 'must-have' 'Acceptable' requirement (matrix ref. 5.07) is a ‘UPS only feed available’ for Open Racks.

D. There is also additional guidance information in the requirements matrix in the Electrical section, for example, the 'must-have’ 'Optimum' requirement (matrix ref. 5.07) that both ‘UPS and non UPS feeds available’ for Open Racks

E. The ‘considerations’ attributes which also have ‘Acceptable’ and ‘Optimum’ parameters, are considered ‘nice to have’, for information only and not required for OCP Ready™ certification,

F. Where there are A), B) and C) parameters within the ‘Acceptable’ and ‘Optimum’ columns for an attribute, these are all selectable options for meeting the requirements.

OCP Ready Requirements Matrix

Linked: OCP Ready Requirements Matrix - 2020A

Ref	Requirements Attribute (Optimum or Acceptable Parameter Required for OCP Ready Certification)	Acceptable Parameter	Optimum Parameter	Notes for Acceptable Parameter	Notes for Optimum Parameter
1	Access				
1.01	Building Access	A) Road level with step and threshold free access B) Road level with threshold (Notes Required)	Loading dock with lift or leveler	An acceptable access design is for the goods-in area to be accessible at road level. If a threshold exists then a temporary ramp, masonite or other suitable boards can be used to span over the threshold.	An optimum access design is to use a loading dock with an integral lift or leveller that allows packaged racks on pallets to be transported directly from inside truck level to the data centre goods-in area. The advantages of this type of loading dock are 1) any truck with or without a lifting platform can be used for deliveries, 2) speed of off-loading and 3) it is a safer method and there is less risk to equipment and personnel.
1.02	Delivery pathway, Loading dock to Goods in	A) ≥ 2.7m (108in) H x ≥ 1.2m (48in) W unobstructed access and threshold free B) ≥ 2.7m (108in) H x ≥ 1.2m (48in) W with threshold (Notes Required) C) ≥2.3m (90in) H x ≥ .9m (36in) W unobstructed access and threshold free D) ≥2.3m (90in) H x ≥ .9m (36in) W with threshold (Notes Required)	≥ 2.7m (109in) H x ≥ 2.4m (96in) W x ≥ 2.4m (96in) D unobstructed access and threshold free	An extra height allowance of .5m (20in) above the height of a 2.21m (87in) Open Rack V2 (Open Rack v2.0) on castors is required because we need to consider that racks are packaged and bolted on pallets and are maneuvered using a pallet truck. If the doorway clearances are <2.7m (108in) H x <1.2m (48in) W but ≥ 2.3m (90in) H x ≥ .9m (36in) W then racks can be uncrated at the loading dock and rolled through to the Goods in Area.	An optimum depth measurement of 2.4m should also be considered if there is a security air lock in place between the external and internal space.

1.03	**Delivery pathway, Goods in to IT Technical Space (White Space)**	**A)≥2.4m (96in) H x ≥1.2m (48in) W unobstructed access and threshold free B) ≥2.4m (96in) H x ≥1.2m (48in) W with threshold (Notes Required) C)≥2.3m (90in) H x ≥.9m (36in) W unobstructed access and threshold free D)≥2.3m (90in) H x ≥.9m (36in) W with threshold (Notes Required)**	**≥2.4m (96in) x ≥1.8m (72in) W unobstructed access and threshold free**	**Width of 1.2m has been calculated as a multiple of a standard width access floor tile of .6m. A doorway clearance height of ≥ 2.3m (90in) is acceptable as the height of an Open Rack v2.0 with castors is 2.21m (87in). The doorway width of ≥ .9m (36in) is acceptable as the width of an Open Rack v2.0 is .6m (24in). If a threshold exists then a temporary ramp, masonite or other suitable board can be used to span over the threshold.**	**Width of 1.8m has been calculated as a multiple of a standard width access floor tile of .6m. As the height of an Open Rack v2.0 with castors is 2.21m (87in) an extra clearance tolerance of .2m (8in) should be sufficient to allow for any variance from this.**
1.04	**Corridor floor rolling load**	**A)≥567 kg (1250 lb) (5.56 kN) B)≥459kg (1012lb) (4.5kN) (Notes Required)**	**≥680 kg (1500 lb) (6.67 kN)**	**Spreader plates with a spreader plate ramp can be used e.g. masonite boards, 9mm (0.5in) thick plywood or hardboard to increase rolling load capacity and protect the floor from damage by the rack castors** ***Reference:*** **A) BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Minimum (medium duty). B) The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992.**	**Spreader plates may not be required.** ***Reference:*** **BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Recommended (heavy duty).**
1.05	**Unboxing/ pre-staging/storag e area floor uniform load**	**≥732kg/m2 (150lb/ft2) (7.17kN/m2)**	**≥1221 kg/m2 (250 lb/ft2) (11.97 kN/m2)**	***Reference:*** **BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Minimum (medium duty).**	***Reference:*** **BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Recommended (heavy duty). The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992. Grade: Heavy & Extra Heavy**
1.06	**Unboxing/ pre-staging/ storage area floor concentrated load**	**A) ≥567kg (1250lb) (5.56kN) B) ≥459kg (1012lb) (4.5kN) (Notes Required)**	**≥680 kg (1500 lb) (6.67 kN)**	**Reference: 1) BICSI 002-2019 – Computer Room Access Floor Performance**	***Reference:*** **BICSI 002-2019 – Computer Room Access Floor Performance Specifications –**

				Specifications – Minimum (medium duty). 2) The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992. 3) The European Standard for Raised Access Floors BS EN 12825:2001	Recommended (heavy duty).
2	Ramps				
2.01	Gradient	1:8-1:12	A) No Ramps Required B) ≤1:12	*Reference* BICSI 002-2019 'The British Industrial Truck Association 16178 (BITA) specifies a maximum slope of 12.5% for forklift trucks, which is a gradient of 1:8 or 7.13 degrees.	A fully populated Open Rack weighing 1500kg (3,300 lbs) would prove very difficult to move up a ramp that was steeper than a 1:12 incline
2.02	Width	≥1.2m (48in)	≥1.5m (60in)		
2.03	Landing area	≥1.2m x 1.2m (48in x 48in)	≥1.5m x 1.5m (60in x 60in)	This does not allow for cabinet rotation.	This does allow for cabinet rotation.
2.04	Railings	<900mm (36in)	≥900mm (36in) and <1000mm (40in)	UK building regulations specify that handrails to be between 900mm and 1000mm above surface or steps pitch line / 900mm and 1100mm above surface of landings.	
3	Goods Lift / Elevator				
3.01	Weight loading	≥500kg (1250lbs)	≥1500kg (3300lbs)	Open Rack v2.0 has a standard 500 kg base load with a lift of this specification could transport (1) Open Rack v2.0 @ 500 kg.	Open Rack v2.0 has a standard 500kg base load and 1400Kg heavy load therefore with a lift of this specification up to (3) Open Rack v2.0 @ 500 kg could be transported together or (1) Open Rack v2.0 @ 1400 kg. 1500kg is a standard maximum lift load specification.
3.02	Door height	≥2.3m (90in) Lift /Elevator door opening height (not internal cabin)	≥2.4m (96in) Lift /Elevator door opening height (not internal cabin)	A door opening clearance height of ≥ 2.3m (90in) is acceptable as the height of an Open Rack v2.0 with castors is 2.21m (87in)	
3.03	Width	≥1.2m (48in) Unobstructed door opening width	≥1.5m (60in) Unobstructed door opening width		
3.04	Depth	≥1.2m (48in) Unobstructed cabin depth	≥1.5m (60in) Unobstructed cabin depth		

4	**IT Technical Space (White Space)**				
4.01	**Floor rolling load**	**A) ≥567kg (1250lb) (5.56kN) B) ≥459kg (1012lb) (4.5kN) (Notes Required)**	**≥680kg (1500lb) (6.67kN)**	**Spreader plates with a spreader plate ramp can be used e.g. 9mm (0.5in) thick plywood or hardboard to increase rolling load capacity and protect the floor from damage by the rack castors.** ***Reference:* A) BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Minimum (medium duty). B) The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992.**	**Spreader plates may not be required.** ***Reference: BICSI 002-2019* – Computer Room Access Floor Performance Specifications – Recommended (heavy duty).**
4.02	**Floor uniform load**	**≥732kg/m2 (150lb/ft2) (7.17kN/m2)**	**≥1221kg/m2 (250lb/ft2) (11.97kN/m2)**	***Reference:* BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Minimum (medium duty).**	***Reference:* BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Recommended (heavy duty). The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992. Grade: Heavy & Extra Heavy**
4.03	**Floor concentrated load**	**A) ≥567kg (1250lb) (5.56kN) B) ≥459kg (1012lb) (4.5kN) (Notes Required)**	**≥680kg (1500lb) (6.67kN)**	***Reference:* A) BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Minimum (medium duty). B) The Property Services Agency Specification for Platform Floors (Raised Access Floors): PSA MOB PF2 PS/SPU March 1992.**	***Reference:* BICSI 002-2019 – Computer Room Access Floor Performance Specifications – Recommended (heavy duty).**
4.04	**Finished floor to ceiling height**	**<3.1m (124in) (Notes Required)**	**≥4.5m (180in),2. ≥3.1m (124in)**		
4.05	**Access floor clearance**	**A) <450mm (18in) if not used for cooling (Notes Required) B) <900mm (36in) if used for cooling (Notes Required)**	**A) No Access Floor B) ≥450mm (18in) if not used for cooling C) ≥900mm (36in) if used for cooling**		

5	**Electrical**				
5.01	**Number of independent circuits to the rack**	**1N (A)**	**2N (A+B)**		**Redundant circuits and paths providing improved resilience**
5.02	**Maximum Circuit Capacity**	**A) 3ϕ 16A/230V B) 3ϕ 30A/208V**	**A) 3ϕ 32A/230V B) 3ϕ 60A/208V C) 3ϕ 50A/208V**	**Circuit capacity requirement for 6.6kw power shelf.**	**Circuit capacity requirement for 12.5kw power shelf. Dual for 25kw power shelf.**
5.03	**Circuit Voltage**	**120 VAC Nominal**	**A) 480/277 VAC nominal B) 400/230 VAC nominal C) 208 VAC Nominal**		
5.04	**Frequency**	**47-63 Hz**			
5.05	**Power receptacle / WIP type**	**A) IEC 60309 516R6W B) NEMA L21-30C or NEMA L15-30C**	**IEC60309 532R6W IEC60309 460R9W Hubbell CS8364C Russellstoll 9R53U2W**		
5.06	**Circuit receptacle location**	**Underfloor**	**Overhead**		
5.07	**Upstream UPS options**	**UPS only feed available**	**UPS and non UPS feeds available**		**There could be a case where the requirement is for Feed A is UPS fed but Feed B is from the utility. Non-UPS (utility) supply to be taken from the input board (or generator / HV distribution board after the ATS) should be considered. See Fig. 8-4**
5.08	**Rack-based batteries permitted**	**Not Allowed (Notes Required)**	**Allowed**		
5.09	**Generator load acceptance time**	**<90 seconds**	**<60 seconds**	**Within hold up time of the BBU at full load**	
6	**Cooling**				
6.01	**Rack airflow direction**	**Front to Back**		**Technology standard**	
6.02	**Air containment methods**	**Hot/Cold aisle containment for all cabinets in white space**	**Hot aisle containment or rack chimney**	**Availability and energy efficiency reasons**	**Majority White Space is cold to enter and to perform maintenance without crossing hot aisles**
6.03	**Maximum rack density**	**≥8 kW**	**≥12 kW**		
6.04	**Minimum cold aisle width**	**≥1200mm (48in)**	**≥1500mm (60in)**	**Air flow capacity.**	
6.05	**Minimum free width cold aisle (inside cage)**	**≥900mm (36in)**	**≥1200mm (48in)**	**Maintenance aisle with rack front door open.**	
6.06	**Minimum hot aisle**	**≥900mm (36in)**	**≥1200mm (48in)**	**Maintenance aisle with rack rear door open. Air flow**	

	width			**capacity.**	
6.07	**Inlet Air Conditions**	**ASHRAE Recommended**	**ASHRAE Class A1 Allowable**	**Energy efficiency**	
6.08	**Air Quality**	**Other (Notes Required)**	**EN 779 G4 and F7 filtering & Gas particulate monitoring to the ANSI/ISA 74.04-1985 G severity levels**		
6.09	**Temperature rise**	**≥8 Deg C DeltaT**	**≥12 Deg C DeltaT**	**Minimum viable**	**Energy efficiency**
6.1	**Cabinet blanking of open space**	**Recommended**	**Mandatory**	**Availability and energy efficiency reasons**	
7	**Cabling**				
7.01	**Cabling infrastructure pathways**	**A) Top and Rear of rack fed B) Bottom and Front of rack fed**	**Top & Front of rack fed**		**Ease of maintenance. Optimum configuration aligns with Open Rack v2.0 power – rear, network – front.**
7.02	**Overhead Network Infrastructure containment levels**	**2 Levels (Intra-Pod cabling; Inter-Pod cabling)**	**3 Levels (Intra-Pod cabling; Inter-Pod cabling; OOB cabling)**		
7.03	**Fibre Type (if installed)**	**A) OS2 & OM3 B) Installed Per Customer Requirements**	**OS2 & OM4**		
7.04	**Fibre connection presentation (if installed)**	**Installed Per Customer Requirements**	**Interchangable LC Duplex and MPO**		

	Considerations Attribute (Optimum or Acceptable Parameter Required for Information - Not Required for OCP Ready Certification)				
	Service				
8.01	**Replacement PSU Modules**	**A) <2 day delivery B) Secure storage available C) Other (Notes Required)**	**In stock onsite**		
8.02	**Replacement BBU Modules**	**A) <2 day delivery B) Secure storage available C) Other (Notes Required)**	**In stock onsite**		
8.03	**Option to monitor PSUs and BBUs**	**No**	**Yes**		**There could be a requirement for the colo facility to monitor the health of the PSU and BBU using its DCIM as part of the SLA**
8.04	**Remote hands for PSU and BBU**	**No**	**Yes**		

	replacement or expansion				
8.05	**Remote hands for OCP IT hardware replacement or expansion**	**No**	**Yes**		
	Efficiency				
9.01	**Site Operations Standards**	**A) Bicsi 009-2019 Data Center Operations and Maintenance Best Practices B) Other (Notes Required)**	**OCP Critical Facility Operations Guidelines**		
9.02	**Site PUE Monitoring**	**A) Periodically measured B) Other (Notes Required)**	**Continuously monitored**		
9.03	**Site Design PUE**	**A) <1.5 B) Other (Notes Required)**	**<1.2**		
9.04	**Site Annualized PUE Current Achievement**	**A) <1.5 B) Other (Notes Required)**	**<1.2**		
9.05	**Site WUE Monitoring**	**A) Periodically measured B) Other (Notes Required)**	**Continuously monitored**		
9.06	**Site CUE Monitoring**	**A) Periodically measured B) Other (Notes Required)**	**Continuously monitored**		
	Openness				
10.01	**PUE Published**	**A) Available upon request B) Other (Notes Required)**	**Real-time available online**		
10.02	**Facility Design Drawings & Files**	**A) Available to view upon request B) Other (Notes Required)**	**Available online**		

Appendix G. Safety and Compliance

It is important to consider the rules and regulations of local authorities that have jurisdiction (AHJ) when planning for and deploying Open Racks, below are some of the considerations that have been identified.

- In the USA ensure UL-VO and FCC certificates are available when racks are delivered to the data center, otherwise there is a possibility that the delivery won't be accepted.

- In the USA CA region and other areas in the world where there is an earthquake risk, ensure earthquake protection i.e. seismic kits are installed in the Open Rack for GR63 Zone 2 and GR63 Zone 4 support.

- If Open Racks are to be installed in a data center with an access floor, ensure that the PDU cables are lengthened as part of the Level 11 manufacturing process, as the cables are normally sized for overhead power receptacle presentation. In the USA onsite modification is not normally allowed.

- It is the responsibility of the owner of the facility, not the manufacturer or installer to identify hazards including arc flash. An arc flash analysis should be performed to help identify risks to personal. This study may be required to be completed or updated with any new electrical device being added to a data center.

  Refer to NFPA 70E and NEC for labeling requirements, established category levels and associated PPE to reduce arc flash exposure.

14. References

- Data Center Facility (DCF)
  http://opencompute.org/projects/data-center-facility/
- OCP Ready™ Facility Recognition Program
  https://www.opencompute.org/projects/ocp-readytm-facility-recognition-program
- Critical Facility Operations - Incubation (CFOPS)
  https://www.opencompute.org/projects/critical-facility-operations-incubation
- Critical Facility Operations Framework
  https://drive.google.com/file/d/1Ry5syU6SvLhWswDa2LflABXHb3XBFnK1/view
- Modular Data Center (MDC)
  https://www.opencompute.org/projects/modular-data-center
- Advanced Cooling Facilities (ACF) - Incubation
  https://www.opencompute.org/projects/advanced-cooling-facilities-incubation
- Rack & Power (R&P)
  http://opencompute.org/projects/rack-and-power/

- How to Become an OCP Colo Solution Provider for Facilities
  https://www.opencompute.org/sp/how-to-become-a-solution-provider

- OCP Marketplace - Circular Economy
  https://www.opencompute.org/circular-economy

- Definition of Manufacturing Levels
  https://www.amax.com/blog/?p=668

- Schneider White Paper - WP 228 - Analysis of Data Center Architectures Supporting Compute Project (OCP)
  https://it-resource.schneider-electric.com/white-papers/wp-228-analysis-of-data-center-architectures-supporting-compute-project-ocp

- Schneider Online Traditional vs OCP capital cost calculator
  https://www.se.com/ww/en/work/solutions/system/s1/data-center-and-network-systems/trade-off-tools/traditional-vs-open-compute-data-center-capital-cost-calculator/

- ANSI/ BICSI 002-2019 - Data Center Design and Implementation Best Practices
  https://www.bicsi.org/standards/available-standards-store/single-purchase/ansi-bicsi-002-2019-data-center-design

- Greenhouse Gas Protocol
  https://ghgprotocol.org/

- ITRenew Data Center Impact Report: The Financial and Sustainability Case for Circularity
  https://www.opencompute.org/documents/itrenew-the-case-for-circularity-data-center-impact-report-ocp-v11-pdf

- Intel - IMPROVE DATA CENTER COOLING FACILITY EFFICIENCY THROUGH PLATFORM POWER TELEMETRY
  https://www.opencompute.org/documents/ocp-wp-dcf-improve-data-center-cooling-facility-efficiency-through-platform-power-telemetryr1-0-final-update-pdf

- The Current State of Data Center Energy Efficiency in Europe - OCP White Paper
  https://www.opencompute.org/documents/the-current-state-of-data-center-energy-efficiency-in-europe-ocp-white-paper

- Data Center Cable Infrastructure Color Schema
  http://files.opencompute.org/oc/public.php?service=files&t=c286ca170f27d7eda91b15a62bad73eb