

A FUTURE AT THE EDGE: EDGE DATA CENTER WORKING GROUP SOLUTIONS BRIEF PAPERS ISSUE 1

JANUARY 2020

Types and Locations of Edge Data Centers: Scoping Locations that Work for your Needs

BY: Hiren Surti (Crown Castle)
Tom Craft (CommScope)

Pack Jones (EdgeConneX)
Tom Widawsky (HDR)

OVERVIEW

With 5G on the horizon, and the promise of a host of new technology applications like autonomous vehicles, console-less gaming and instant grocery delivery, we are set to create an unprecedented amount of data that will require rethinking the way our telecommunications and data infrastructure is built and managed going forward.

Users now expect much more functionality from service providers to support IoT devices, virtual reality applications, connected vehicles, telemedicine, immersive media, and other products and services that require faster speeds, higher capacity and larger throughput on-demand. The demands of these new technologies – instantaneous processing and transmission of large amounts of data – require new methodologies for data management.

To meet the needs of new applications, data needs to be hosted very close to users and network functionality needs to be performed in milliseconds. One of the most promising solutions is managing data at the edge of the network. What does that mean? Right now, much of our data is processed and stored in centralized data centers located sometimes many cities away from an end user. Moving away from these traditional large data centers located far away to small data centers at the “edge” of the network – like potentially a city streetlamp – will help keep data close to the user and fulfill the objectives of latency, bandwidth and quality of service (QoS).

Bell Labs predicts that 60 percent of servers will be placed in an Edge Data Center (EDC) by 2025.¹ With this significant shift on the horizon, we need to work across the industry, with our partners, to account for connectivity, latency, bandwidth, redundancy and reliability. We need to create common terminology, requirements and standards for EDCs.

¹The Future X Network: A Bell Labs Perspective by Marcus K. Weldon (2016-03-01), CRC Press

This paper reviews and establishes edge equipment as well as types and potential locations of EDCs. We focus on potential locations for two types of edge data centers: network-based and establishment-based. There is a potential third type of EDC – device-based EDCs – that are outside the scope of this current paper. This paper is the part of a series of “Solutions Briefs” that will set a framework for the future of EDC development and industry standards.

We welcome comments, feedback and new ideas and we take the next steps to prepare for a future at the edge. TIA and the members of the Edge Data Working Group invite others in this ecosystem to join in the next steps of development of these concepts. Contact EDCINFO@tiaonline.org for more information.

I. EDGE EQUIPMENT LINEUP

The edge compute function will sit within a lineup of single to multiple cabinets, or racks. It will aggregate with the radio access network (RAN) to offer user access from wireless sites. An example of such a lineup is shown in Figure 1. This configuration would be used inside a Central Office (CO) or core data center.

Obviously, for larger EDC sites the equipment count or function would scale up. Likewise, for smaller EDC sites (for example, single rack deployments) the equipment count or function would scale down. Although a figure is not shown for a smaller system, rather than employing optical distribution frames (ODFs) a fiber shelf will suffice. The same would be true for switches, compute, storage, RAN, and power equipment.

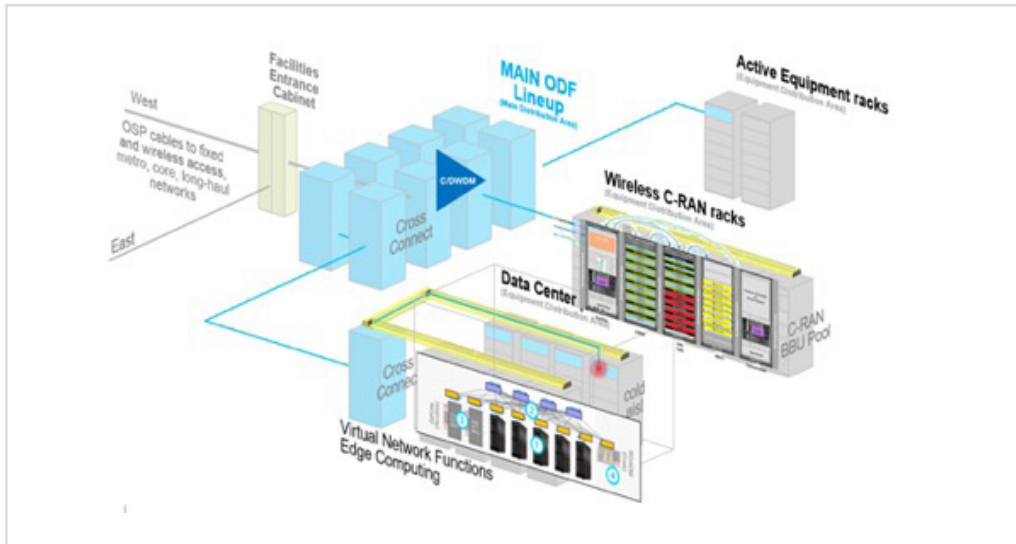


Figure 1. Typical Equipment Lineup for an Edge Datacenter Facility with Legacy Telecom and C-RAN Gear (Courtesy of CommScope)

The estimated size of EDCs as defined by the Central Office Re-architected as a Data center (CORD) architecture is:

- Large: 12 to 49 racks with a power draw of 12 to 14 kW/rack (or cabinet).
- Medium: 4 to 12 racks with a power draw of 8 to 12 kW/rack
- Small: 1 to 8 racks with a power draw of 6 to 8 kW/rack

The function of an EDC from ODF, network switches, transport, compute, and storage hardware perspectives are also out of the scope of this document. However, it is assumed that the necessary equipment for one, and up to 49 racks of equipment, will scale accordingly and perform the same function. This assumption should hold true, regardless of whether the EDC is deployed in a single OSP cabinet, at a radio site, or in a multi-rack lineup within a core site.

II. EDC LOCATIONS

Locations overview

At a high level, the edge computing locations divide into three types of locations, network-based sites, establishment-based sites and device-based sites. This paper focuses only on network-based and establishment-based sites and looks at classifications within these categories. The device-based edge is beyond the scope of this paper but will be addressed in the future.

1. Network-based EDC Sites

- ✓ Cell/radio site
 - Micro-modular DCs
- ✓ Access sites
 - Utility/Industrial Internet of Things aggregation point
 - Micro-modular DCs
 - Other access sites
- ✓ Aggregation sites
 - Regional DCs
 - Colocation DCs
 - Interconnection point-of-presence (POP) (not as part of an Internet exchange point (IXP))
 - Other aggregation sites
- ✓ Core or central cloud sites
 - COs
 - Core DCs
 - IXPs
 - Other core sites

2. Establishment-based EDC Sites

- ✓ On/off premise locations at a venue
- ✓ Enterprise sites
- ✓ Industrial sites
- ✓ Retail sites
- ✓ Financial sites
- ✓ Railway (or other transportation related) sites

3. Device-based EDC sites

(Not covered in this paper)

Network-based EDCs

The network-based edge results in a tier-based architecture and is uniquely structured to support EDCs. This approach is favored because a single EDC existing in a network will have an elevated possibility of failures and compromise the availability of the network. As shown in Figure 2, a tiered edge computing site, or other network-based location, guarantees high availability. This is true not only for mobile networks, but also for the associated information and communications technology (ICT) networks.

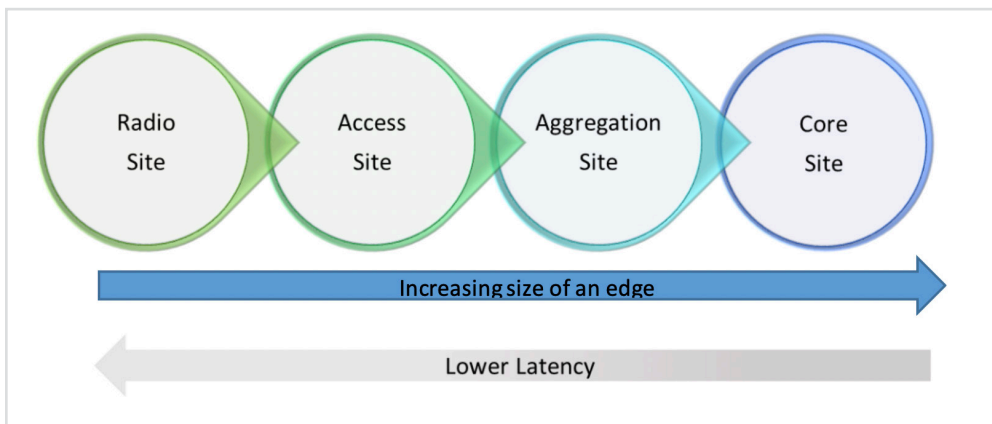


Figure 2. Network-based edge tiered locations

1. Radio and access sites

- Radio site: As latency is an important element of edge computing, the radio site is typically located at the cell tower and small cell locations. The priority is to have them one hop away for both mobile and non-mobile users to achieve ultra-low latency. These locations are typically unmanned and consist of outside plant (OSP) equipment. However, these radio sites process much of today's mobile Internet data since they are collocated in high-service areas.
- Access site: This system typically terminates much of the fiber and offers robust network flexibility points as well as suitable locations for deploying EDCs. However, there is considerable variability in these locations. Therefore, proportional changes in latency can be expected. The advantage is that they provide another resilient tier of potential locations for edge computing. These sites are used for data collection, as well as short-term and immediate data processing needs.

Types of radio and access EDC sites

These radio and access EDCs can range from 1 to 8 racks with a power draw of 6 to 8 kW/rack. Their power availability varies from 50 kW to 150 kW in capacity. However, the power availability can be greater depending on the ground space and tenants available at the site.

- Small rack count systems are deployed on a concrete pad in an outdoor cabinet or closure.
- For multiple racks, (4 to 8 in this case), a larger walk-in EDC is deployed. These are typically seen as containerized or purpose-built large cabinet/shelter types. They can be used for single or multitenant purposes.
- In cases where the radio, or access site, has a hut, or shelter, self-contained indoor racks can be deployed within the shelter. If it is self-contained, it would have its own power distribution, fire protection/suppression, and close coupled cooling. Examples of these infrastructure components are shown in Figure 3.




Figure 3. Radio and Access Edge Datacenter: A) Single Rack Cabinet (Courtesy of Dell); B) Multi-Rack Walk-in Cabinet; C) 2 Rack Self-Contained Indoor System (Courtesy of CommScope); D) Large Containerized System (Courtesy of Vapor I/O)

2. Aggregation sites

Aggregation edge systems are comprised, typically, of multiple access networks. They are connected, via fiber, at a point prior to traffic being routed to the regional data centers, or points of presence (POPs). The local aggregation sites are where switching and routing equipment aggregates access traffic and are suitable locations for deploying edge computing hardware. They can also include meet-me-rooms that are ideal carrier handover points. The primary use case will be for user data collection, data aggregation, and analytic processing and transmission.

Types of aggregation site EDCs

Typically, aggregation sites have buildings for the installation of indoor rack line-ups as defined by Telecommunications Industry Association (TIA), Building Industry Consulting Service International (BICSI), and other data center-centric standards. The size of these data centers is medium to large and consist of rack counts from 4 to 49 racks, with a power density of 8 to 14 kW/rack.

In some cases, these sites will require the deployment of containerized, or purpose-built, prefabricated EDCs. The size of these particular systems will also be medium to large, with similar rack and power configurations.

3. Core sites

Core sites, generally, are centralized and contain the core network equipment. These sites orchestrate the needs between access and aggregation sites and perform hyper local data collection or short-term/immediate processing, separate from standard core data processing and storage.

Since these sites have the necessary infrastructure, they tend to have higher availability ratings. At the same time, they are usually in close proximity to people. These factors make them high-value EDC locations. The tradeoff of edge access equipment/function and latency for users will need to be evaluated for core sites. Latency at these locations tends to be higher than at the radio, access, and local aggregation sites. However, they tend to be manned locations and are easily accessible.

Types of core site EDCs

In general, core sites have buildings for the installation of indoor rack line-ups as defined by TIA, BICSI, and other data center-centric standards. The size of these data centers is large and they consist of rack counts from 12 to 49 racks with a power density of 12 to 14 kW/rack.

In some cases, these sites will require the deployment of a containerized or purpose-built prefabricated EDC. The size of these particular systems will also be medium to large, with similar rack and power configurations (see Figure 4).




Figure 4. Examples of indoor rack line-ups or pods to suit the in-building deployments. Pictures are of APC Schneider Electric's EcoAisle™ product line (left). This lineup is built up from a series of single rack or cabinet products with the appropriate aisle containment infrastructure, power distribution, and possibly in-row coolers. Another indoor deployment approach is Vertiv's SmartRow™ system (right). This system is smaller in size and is built from a series of self-contained cabinets with in-row coolers,

Purpose-built modular data centers, for outdoor deployments, are shown in Figure 5. These can be deployed in core DC yards and parking lots to industrial pad sites to support the edge functions. The first example is from Dell (MDC) and is shown with its power and cooling plant. The second, a smaller outdoor system from Vertiv called SmartMod™, supports smaller rack counts, with cooling and power distribution.



Figure 5. Large Modular Datacenter Facility (Courtesy of Dell); SmartMod™ Large Walk-in Closure (Courtesy of Vertiv)*

Establishment-based EDCs

Establishment-based edge computing certainly can leverage the tiered architectures discussed previously, for off-premises configurations. On-premises EDCs take advantage of collocation space, small on-site data closets, or server rooms. Since the EDCs are collocated within these facilities, or at nearby locations, and placed with respect to service demand, latency tends to be on the lower end.

The establishment-based EDC is, typically, private-network driven. It orchestrates transport needs between access and aggregation sites and performs hyper local data collection or short-term/immediate processing separate from standard core data processing and storage.

Types of establishment-based EDCs

- Indoor racks as defined by TIA, BICSI, and other data center-centric standards.
- Rack counts are from 1 to 8 racks with a power draw between 6 to 8 kW/rack.
- The racks tend to be a self-contained cabinet with close coupled cooling. This is because typical data closets, and enterprise data centers, are not equipped to handle the power distribution and cooling requirements.
- In some cases, where indoor deployments are not practical, OSP cabinets, or purpose-built EDCs, are a good fit.

Examples of these systems are shown in Figure 6. The first two examples show a modular system from RakWorx, in one-rack and three-rack options. These systems can be sized to meet edge compute needs.

The last enclosure is a small shelter from Solarcraft designed to house and protect a few racks of active gear. It can be deployed, for example, in the yard of most enterprises, schools, or hospitals to support the EDC when indoor space is not cost effective.

III. CLOSING THOUGHTS

What, and ultimately where is “the edge?” is a question that has been the topic of discussion in many board rooms, labs, and white papers. While there isn’t one answer, we do know that the edge function will be deployed in data center-like infrastructure positioned close to its users. Edge Data Centers will be a reliable low-latency link with high bandwidth for potentially massive amounts of data and will provide an unparalleled experience for said service level agreements (SLAs).

Use case demands will require EDCs to be located in a tiered architecture, as defined in this document, in facilities that are ready for such active equipment. All EDCs will need to be within reasonable proximity of the end user. Many will be in the closets of small enterprises that want to host local compute/storage for security reasons, yet have lack of real time access, over the same transport, to the cloud.



Figure 6: Establishment-based EDCs; A) 1 and 3 Rack Edge Easy™ DC (Courtesy of RakWorX); B) UltraStrut™ Outdoor Shelter (Courtesy of Solarcraft)

Depending on user needs, there will be many deployment locations and options. One would be to install the EDCs in a small rack count, outdoor cabinet. Other options range from an indoor rack to a large purpose-built outdoor container designed to house and protect the edge function equipment.

Scoping potential EDC site and location types is just the beginning. Many other topics need to be addressed to fully define EDC systems – including power, cooling, security, reliability/availability/serviceability (RAS), performance ratings, and operations/automation/maintenance/administration (OAMA). The “A Future at the Edge” solutions briefing paper series will address these areas in future papers.

To learn more about the TIA’s Working Group efforts in developing information and standards on Edge Data Centers (EDCs), go to the TIA website: www.tiaonline.org. If you have opinions or expertise to lend to this effort, please reach out to edcinfo@tiaonline.org.

* Disclaimer: The information and views contained in this article are solely those of its authors and do not reflect the consensus opinion of TIA members or of TIA Engineering Committee TR-42. This article is for information purposes only and it is intended to generate opinion and feedback so that the authors and TIA members can learn, refine, and update this article over time. The Telecommunications Industry Association does not endorse or promote any product, service, company or service provider. Photos and products used as examples in this paper are solely for information purposes and do not constitute an endorsement by TIA.