What is a Rack Unit?

It seems straight forward. A Rack Unit is 1 ¾ inches of mounting space in an equipment rack. Either the space has equipment mounted in it currently, or the space is free capacity. Or is it?

The lowly Rack Unit has more to it than that. It has a relationship to power, network connectivity, floor space, heat dissipation, rack type, the IT technology Stack, applications, and data center policies. This makes it a more complicated unit of capacity than it first appears.

New data centers are designed for 20 years of capacity usage. This means an enterprise makes an informed investment to commit a certain amount of capital resources with the expectation that this investment will support the revenue generating capability of the enterprise for 20 years.

Of course, no one knows what technology will be used during the entire 20-year life span! Data center architects must use their best guess, along with proven systems of analysis, to create a data center with enough flexibility and capacity to meet the goals of the enterprise.

From an executive's perspective, that means capacity to support revenue generation for 20 years.

Rack Units are units of capacity but they are intertwined in a complex relationship with many other capacity units with a single, final goal … support revenue generation.

This paper will explore the relationships between various data center metrics including Rack Units. It will use two methods of analysis: Aggregation and Integration. We will keep it simple and focused. This is a practical analysis, not a theoretical treatise.

Revenue Generation, a Service Viewpoint

Revenue is received in exchange for services or products. This may seem far removed from Rack Units but there is a path between the two and the relationship will pop out as we follow this very important path.

Just as a path in a forest is the aggregation of numerous small steps, the path from Rack Units to revenue is an aggregation of numerous small items with a common property. We call them … Service Units.

We will use Service Units to show the common, basic relationship between many metrics and our lowly Rack Unit. This understanding can be a very powerful tool in your analysis, planning and usage of data center resources.

Since we are considering Service Units as a base or core thing that can be aggregated, let's explore a simple yet extreme example of how and why this works. If the concept holds up in this case, it is probably usable as our core concept.

Company XYZ sells physical products … drill bits. The company generates revenue by exchanging its drill bits for money. How does this relate to services?

Why do consumers buy drill bits? They need holes! That's what they really need. They buy drill bits as a means to get what they really need. If someone could create a service that would create a hole of the exact size and in the exact place people required for a price comparable to buying drills, electricity, drill bits and the labor to use them, then no one would make or buy drill bits as a product.

A drill bit, or any physical product, can be thought of not only as a service replacement due to cost or technological limits, but also as the aggregation of all the products and services used to create and deliver it to the consumer.

Once we change our viewpoint of a physical item as being a thing to being a replacement for services that either would cost more or are technologically impractical, it is easy to see how aggregation occurs and can be measured.

This is not academic. It is core to the understanding of what a Rack Unit is. An equipment rack provides a service. It houses equipment. What do we mean by housing? The rack provides everything the equipment needs to perform **the service the equipment is designed for**. The rack provides a safe, secure specific location for the equipment. Because the equipment has a specific location, it can be reliably powered, cooled and connected to for data transfer.

The equipment can join into the wonderful interplay of artifacts and processes that result in a final service or product (aggregation of service) delivered to a consumer … because the equipment has been housed by a rack.

The rack has a certain capacity to perform this service. And a Rack Unit is one measurement of the rack's capacity to perform a housing service. This housing service joins into the aggregation of many other services toward the ultimate goal of generating revenue by serving consumers.

What is a Service?

The word "service" has cropped up a lot in this document. A better explanation of it is required. A service has certain common properties regardless of its type. A service performs a **function**. It has a certain **capacity** and **availability** to perform the function. During aggregation, the function of a service is **transformed** into a supporting element of another service.

Understanding this transformation is the key to understanding how the aggregation can be measured. We will use our lowly Rack Unit to help in this understanding.

A rack performs the function of housing equipment. This is a service because, though it is economically and logistically impractical, it is possible to hire a person to do nothing but stand in a spot and hold a piece of equipment, giving it a secure and specific place where power and data connectivity can be supplied. If a person was hired to perform this silly function, they wouldn't be considered to be making anything. They would be providing a service similar to mowing a lawn or holding a sign at a street corner while wearing a silly costume to attract attention.

Obviously, it is less costly and more practical to buy a rack to perform the function of housing equipment, and a rack has a certain capacity to do so. We use Rack Units to measure the theoretical capacity of a rack but the actual capacity depends on whether any specific Rack Unit is actually **available for an intended purpose**.

The availability of a Rack Unit depends on the properties of the item requiring housing as a factor in its ability to perform its function.

This important relationship is the basis for understanding **transformation** of function. When a planner decides to use a set of rack space to house a server, he/she is making a decision to transform the available housing capacity into another functional capacity … application hosting capacity. Even before the actual server is purchased and installed, the

decision to transform available housing capacity into application hosting capacity is made by the planner.

When the server is mounted in the rack, part of the housing capability of the rack is transformed into a **supporting element** of the service or function performed by the server. The service performed by the rack has been aggregated into the service performed by the server. When an application is installed on the server, part of the application hosting capacity of the server is transformed into the function provided by the application. When the application is used to provide a business service to a consumer in exchange for revenue, the aggregation from Rack Unit to revenue is complete.

Stranded Rack Units

A Rack Unit is a measure of capacity to perform the function of "housing" only when the Rack Unit is **actually available for an intended purpose**. If we want to transform a unit of housing capacity into a unit of application hosting capacity, as in the above scenario, the application hosting function requires more than just housing to support it. It requires power, cooling, network connectivity and the available of contiguous housing capacity becomes an issue when the equipment performing the application hosting is larger than a single Rack Unit.

Stranded Rack Units is a measure of **non-availability** for a purpose.

There have been many proposed methods for calculating Stranded Rack Units based not on any specific purpose but on the general usage or expected usage of the entire data center. As an example, it can be calculated that the average power draw per used Rack Unit is 0.1 Kwatts and so, on average, future capacity requirements will probably average the same power draw.

Likewise, the average rack space used per piece of equipment might be 1.8 Rack Units and, on average, future equipment might also be the same size.

But averaging can get us into trouble. The average temperature for both summer and winter might be a comfortable number, therefore no need for winter coats or summer air-conditioning!

Using averaging in metric calculations has value in overall capacity forecasting but often falls short when applied to specific capacity requirements for real projects that pop up. They demand specific functional capacity as a supporting element. How do we know if this particular capacity is available?

Integration

Integration is the capability of a knowledge system to bring data in from multiple sources, resolve the differences and discrepancies between the sources and form a common, usable pool of information. It is also the ability to relate the basic requirements of a query to the pool so that the answer provides useful and relevant information to accomplish the goal of the query.

Integration is performed via logic expressed as programmatic code or rules. But keeping the code from turning into an unmanageable mass of complexity requires a very intelligent data model as a base to work upon.

The basis of any good data model is common elements that can be easily related across multiple contextual viewpoints. This is where our focus on Service Units finally fits into the message of this paper.

We saw earlier that using the viewpoint of Service Units helped to resolve the apparent difference between a physical product such as a rack and a pure service performed as a substitute for the product.

Understanding the deeper nature of Service Units enables the resolution of many other relationships not immediately apparent to common thought.

If we consider the nature of a Service Unit, it has 5 basic properties:

- Function performed
- Form Factor
- Availability
- Capacity
- How it aggregates with other Service Units

Integration is where the rubber meets the road. Many thousands of person days have gone into the development of the OBTAIN data model. Customers are beginning to recognize that documenting floor space, racks, equipment, power and connectivity via spreadsheets and Visio diagrams, or tools that create little more than pretty pictures provides very little payback when the big questions are asked.

Questions concerning capacity, forecasting and cost require a highly integrated data pool to enable aggregation. Please join in our discussion groups. The OBTAIN developers and support people care very much about questions such as “What is a Rack Unit?”