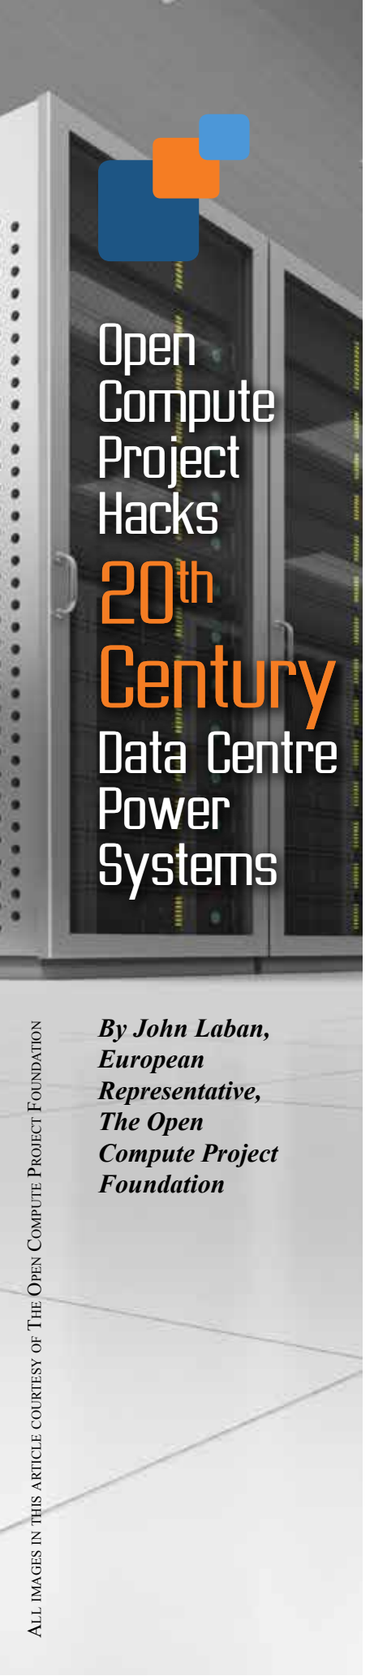

Open Compute Project Hacks 20th Century Data Centre Power Systems

By John Laban, European Representative, The Open Compute Project Foundation

Open Compute Project Foundation

Patents and open source

The not for profit Open Compute Project Foundation was formed in 2011 after several of the world's largest data centre players got together and decided to open source their designs for hardware: servers, storage, network gear and the wider Data Centre eco systems.

The Open Compute Project (OCP) is reimagining hardware, making it more efficient, flexible, and scalable. The OCP community of technology innovators and leaders working together are breaking open the black box of proprietary IT infrastructure to achieve greater choice, customisation, and cost savings.

Since 2011 the OCP open source hardware community has been disrupting the @scale century Data Centre industry and today many of the world's largest data centres are filled with millions of OCP servers.

In 2016 the world`s Telcos formed the OCP Telco Project and they are now transforming their Telephone Exchanges (Central Offices) into OCP virtualised data centres. CORD (Central Office Re-architected as Datacentre) transformations are happening worldwide and things are moving very fast to complete the majority of transformation before the end of this decade. In effect the Telcos are copying the technology solutions of the OTT (Over-The-Top) players like Facebook, Microsoft, Google and Apple that are cannibalising the Telcos revenues.

From the very beginnings OCP gear started to permeate into the large North America banks like Goldman Sachs and Fidelity. 90% of the servers Goldman Sachs buys today are OCP servers.

It's an interesting exercise to study the share price trends during the past two years of the traditional proprietary hardware vendors and the OCP open source hardware vendors and as is often the case with disruptive technologies the financial analysts see the trends at least one year before the early majority adopters.

To help set the scene on why openness always wins in the end I would like to bring you back to open source software. The OCP Foundation is effectively doing for open source hardware what the Linux Foundation (started in 2000) has done for open source software. Today Linux is used on over 90% of the world's servers and what's interesting is that OCP hardware adoption seems to be following a similar growth trend. IDC predicts 50% OCP adoption by 2020 and if this trend continues then we should comfortably achieve 80% adoption rates for OCP open source hardware by 2025. I would recommend the book "The Cathedral and the Bazaar: Musings on Linux and open source by an accidental revolutionary" by Eric Raymond which makes a very convincing case that open source collaborative communities produce better products faster. Perhaps not having patents is part of the success of the Tesla electric car business and the construction of the world's largest machine "The Internet". To help with the global macro economics of what`s happening read "The Zero Marginal Cost Society" by Jeremy Rifkin where he covers the rapidly accelerating worldwide trend towards open source collaborative commons in many industries including distributed power generation.

Killing your babies in 20th Century Mission Critical Data Centres

You are probably thinking what the F*ck has the above got to do with mission critical data centre power systems? Well it's about how open source collaborative communities disrupt the status quo and this relates to people's mindsets because everything we use in data centres is designed by humans. Let me now jump straight into some of the disruptive innovations by telling you what OCP data centre eco system kills in a traditional 20th century data centre.

OCP Data Centres kill:

- Centralised UPS and Batteries and those rooms that house them
- Access Floors
- Intelligent AC Power Strips in the rear of server cabinets
- Tools
- AC/DC Power Supply Units in servers
- 20th Century Stupidity
- Complexity
- Hardware redundancy
- Support Technicians
- Money, energy and material waste
- Low quality – open source produces better products
- Proprietary hardware and user lock-in

Now for those experienced data centre professionals reading about the above deaths (especially when you believe what I believe that it`s true!) you will feel as though I have killed your baby and this means you will experience emotions. I have shown below the Kubler-Ross Change Curve to help you understand where you are as you move into the new disruptive fast-paced open source world of OCP data centres.

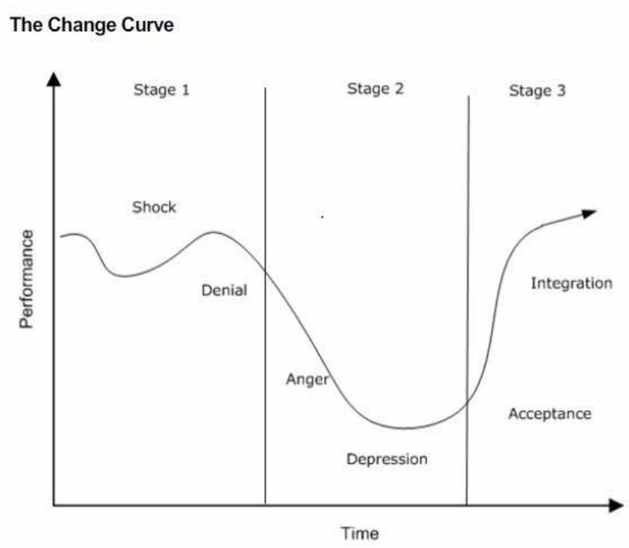

I had been tracking OCP gear from day one but my big eureka moment was when I was immersed in the radically different OCP gear during some consultancy work at the Facebook OCP data centre in Lulea Sweden. Yes I did experience shock that this was so much better than what I had been doing especially since I had been involved in data centres from the earliest days and 25+ years ago as the BT Design Manager for London's new Docklands Development I worked on the design team for Telehouse One (LINX One) in London's Docklands and as we all now know which became the first seed that sprouted the huge data centre cluster we find there today.

Since my eureka moment at Lulea I have immersed myself in the OCP open source community and am now having fun playing with the best people in the data centre industry. I even wrote a one day OCP Awareness Course in December 2015 …in Jan 2016 when the Telcos joined OCP and set up the Telco Project I rewrote the course and now it`s two days. Yes, things move fast in open source communities!

Hacking the Data Centre Electrical Subsystem

The following 3D diagram shows a typical traditional data centre's power and cooling subsystems.

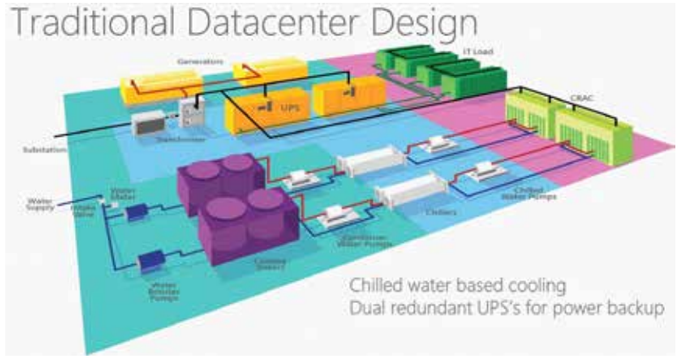

The following 3D diagram shows an OCP optimised data centre. Note that the dual centralised UPS and the mechanical chillers shown in the above 3D diagram have been removed. The LES (Local Energy Store) consists of Li-ion BBU (Battery Backup Units) which are integrated with the N+1 rectifier shelves located inside the OCP v2 Open Racks.

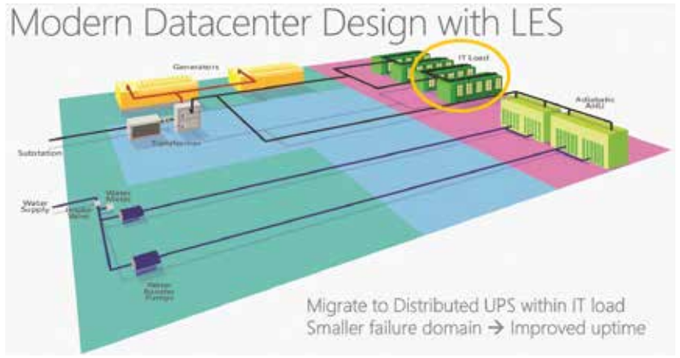

The above 3D diagrams were provided by Kushagra Vaid (GM Microsoft Azure Data Centre Hardware Infrastructure).

For each Microsoft Data Centre the removal of these items result in capital cost savings on each of their large data centres of tens of millions of dollars. This design also significantly improves the data centre PUE (Power Usage Effectiveness) which for the new Microsoft Azure Data Centres is better than PUE=1.125. Similar energy efficiency is achieved for other OCP Data Centres, for example Facebook achieve 1.08 and Facebook show their PUEs at *https://www.facebook.com/PrinevilleDataCenter/app/399244020173259*

I must also add that OCP vanity free open source servers are more energy efficient than traditional servers. Detailed analysis by CERN who use OCP for their HPC (High Performance Computer) show OCP servers are almost 30% more energy efficient than traditional servers.
As you will see in the above 3D diagrams the OCP Data Centre is far simpler than the traditional. The pie chart below from Capitoline www.capitoline.org shows failure mechanisms in traditional Data Centres. OCP Data Centre power systems are more reliable and the failure of a rectifier shelf or the Li-ion BBU only impacts a single rack unlike the large potential blast area (the complete Data Centre) when there are centralised UPS failures.

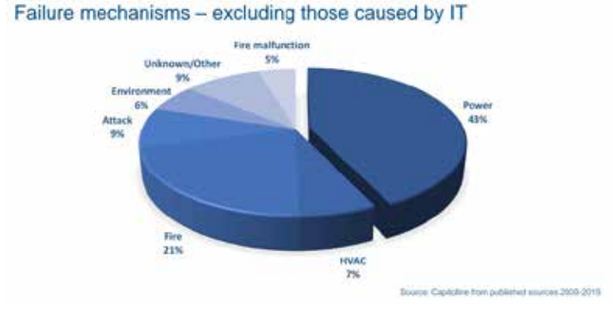

Understanding the makeup of an OCP server

Below is a photo of an OCP vanity free open source server. Note that this server does not have any AC/DC PSU (Power Supply Units). Traditional servers would have dual PSUs running at less than 40% utilisation each and resulting in energy efficiencies of 80%. OCP servers are fed from an N+1 rack level rectifier power shelf which provide significantly higher energy efficiencies. The new OCP rectifier shelf solution submitted by Google in 2016 achieves energy efficiencies figures better than 97.5%

It takes its power via a 12V DC busbar in the rear of the OCP v2 Open Rack. The two relatively large 80mm dia server fans on the rear are considerably more efficient than the cluster of six or eight 40mm fans used in traditional EIA one U servers. (See Fan Cube Law). These OCP servers are designed to minimise heat shadowing and the taller profile supports the larger more energy efficient 80mm fans and allows for much larger heat sinks on this two socket CPU server, and all these factors combined permit higher intake temperatures and this eliminates the need for mechanical chillers. These OCP servers are sometimes referred to as "ambient servers". There are often no VGA graphic elements in these OCP servers and then again this is another reason for lower energy and lower cost. As with all OCP gear it is completely tool-less and the green plastic points are what's called "Green Touch Points" which as the name implies is what you touch to remove/add a component in the OCP server and yet another reason why OCP data centres can achieve technician to server node ratios of 1:25,000 which results in a technician to server node efficiency ratio that is at least one hundred times better than in a traditional enterprise data centre.

Should you wish to know more about this particular OCP server then this short 3 minute Explainer YouTube Video short will help. *https://www.youtube.com/watch?v=tHer0URWUQI*

OCP v2 Open Rack

Graphic below shows an OCP v2 Open Rack.
I should also point out that an OCP Open U is slightly taller than a 19inch relay rack EIA U.
The rear of the rectifier power shelf connects to the 12V DC busbar in the rear of the rack.
The pitch between the vertical rails is also slightly wider as this increases the space/volume efficiency inside the rack. This OCP width also allows for five 4inch wide HDDs and this allows the OCP rack to support 25% more hard drives than a traditional rack. The external dimensions of this OCP rack are no larger than a traditional rack. The version of OCP rack can support loads up to 1,450 kg and they are also slightly taller than an average traditional rack to facilitate the power shelf and Li-ion BBU.

Below in the graphic "18OU" is 18 Open U's and not 180 U's!!

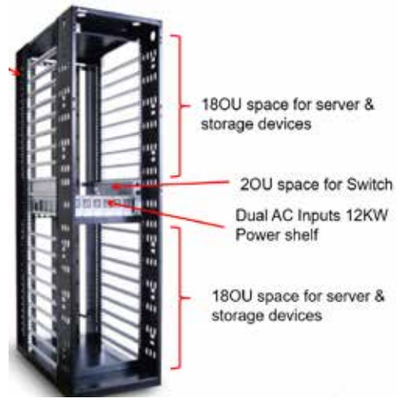

There is also a short OCP Explainer video here for the power distribution in the OCP v2 Open Rack. *https://www.youtube.com/watch?v=mCUgMxrr3nM&t=12s*

OCP data centres strive for simplicity. e.g. Facebook data centres worldwide house tens of thousands of OCP racks and millions of OCP servers yet only has seven different rack configurations.
I hope this provides a good starting point in understanding more about OCP data centres and the OCP open source gear. Here to help if you would like to know more.

This Youtube link of John Laban's presentation in New Zealand provides us with a glimpse of key elements of OCP including: energy efficiency, simplicity and vanity servers: *https://www.youtube.com/watch?v=LFtPK_l4t60*

John Laban
European Representative,
The Open Compute Project Foundation
+ 44 7710 124487
john.laban@opencompute.org