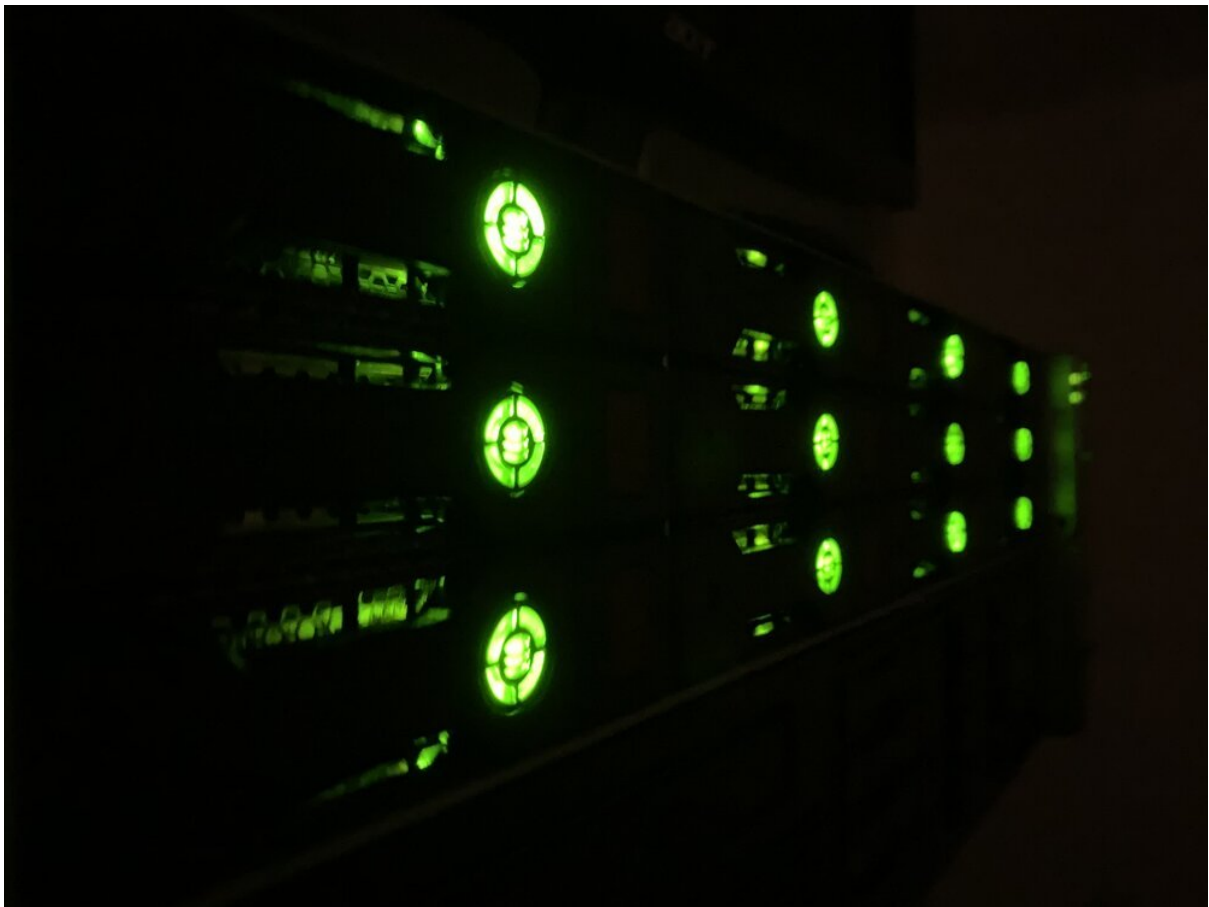# Building a professional 72TB NAS

Configuration and installation of an HPE ProLiant DL380 server

Willem L. Middelkoop
Nov. 22, 2019

Over the past few years I have been moving my data and work from local computers (mostly laptops) into the cloud. Cloud computing is done by servers in a datacenter, powerful computers that do the hard work. As my company grew, I needed more capacity. It was time to add some power to my cloud!

Cloud Computing

Over the past few years "The Cloud" has been marketed as a magical phenomenon solving many business problems. From a technical perspective however, it is nothing more than concentrating computer power in datacenters instead of on your desk.

The main advantages are that you can access (extremely) powerful computing resources that scale as your work load grows. Imagine your desk: there is only so much space for a computer. In contrast, a datacenter is physically much larger and allows many different server computers to be working in concert, offering more capacity.


Servers in a datacenter: the black racks are like closets for the grey server computers, each rack fits multiple servers (photo: HPE).

In addition to scalability, the datacenter is a much better environment for computers than your office. Think of simple thinks like dust, spilling coffee, power outages and differences in temperature and humidity. And if you're like me, commuting by bike, it's a good thing that your data is safely stored in a datacenter instead of on a laptop in your backpack (exposed to shocks, weather, etc).

HPE DL380 Server

My company uses HPE servers as they have a proven track record for reliability and offer many smart design features that allow easy maintenance. In many ways they are designed to make a system engineer's life easier.
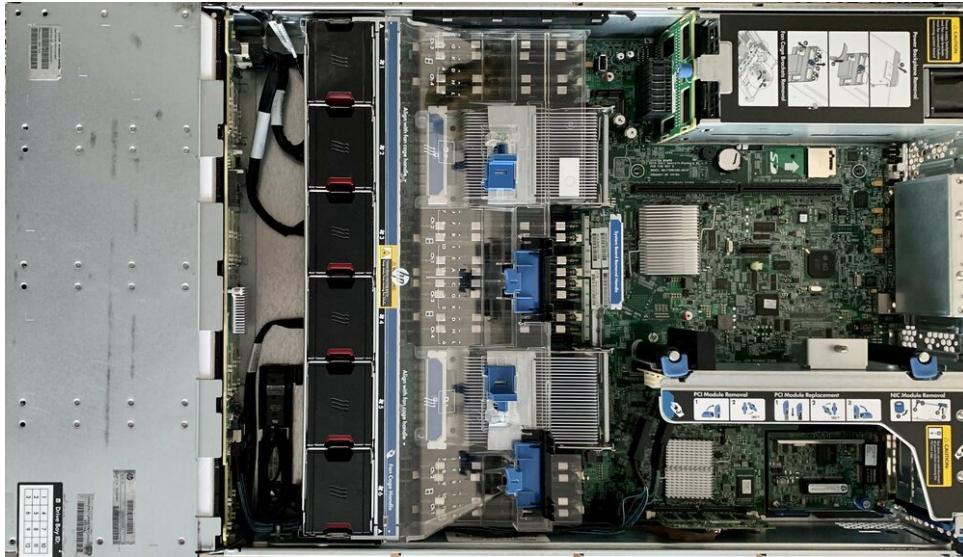
Inside a HPE ProLiant DL380p Gen8 server (LFF)

On the picture above you'll see the HPE DL380 server on the inside. Let me explain it's innards from **left to right**:

- **Drive bay**: the grey metal enclosure is where the harddisks or SSD's live. They are on the front of the server when it's mounted inside a rack, allowing easy access to add/replace drives; usually without requiring any tools.
- **Fan assembly**: the 6 black boxes are fans, sucking cool air from the left and pushing it to the right. Each fan is individually monitored. If a fan fails, others will take over it's load by blowing faster. Each fan is very easy to replace, again without any tools.
- **Processors and memory**: In the center, behind the transparent plastic (cooling tunnel) are the two processors and 16 memory banks. You can recognise the processors by their metal heat sink. They are surrounded with (many) memory banks, physically positioning them very close to the processors allows for very high bandwidths. They are right next to the fan assembly, maximising the thermal envelope.
- **Power Supply Units**: On the top right there are the two power supply inits. They work in concert, dividing the electric load between them. If one fails, the other is capable of supplying the entire server. This provides redundancy as it allows you to connect the server from multiple (different) power sources in de datacenter.
- **System board**: the green chips you'll see on the center right comprise the main system or logic board. It contains various components that make up the server. You'll notice thinks like the (clock) battery, SD-card reader, and iLO-chip. The latter is HPE's integrated 'lights out' system, a computer inside a computer, allowing engineers to remotely check the server's innards.
- **RAID Controller**: On the lower right is the RAID controller, where RAID stands for "redundant array of independent disks". It connects and organises the physical drives in the drive bay. It is capable of dividing data over different disks, allowing for some redundancy in case of drive failure.

Installing an operating system on the server

Just like your desktop or laptop computer, the server needs an operating system too. For servers it is common to use a special kind of operating system that allows for other operating systems to be installed remotely. The main operating system allows you to divide the capacity into multiple virtual servers, this is called an hypervisor OS.

The advantage of using multiple virtual servers is that you can more easily transfer work loads between physical machines without the need to reinstall software. The virtual servers are generic and capable of running on different types of hardware. This enables scalability and redundancy if you use multiple physical servers in combination with multiple virtual servers.

HPE ProLiant DL380p server powered on, green lights indicate the drives are working, monitor is connected and showing the boot process
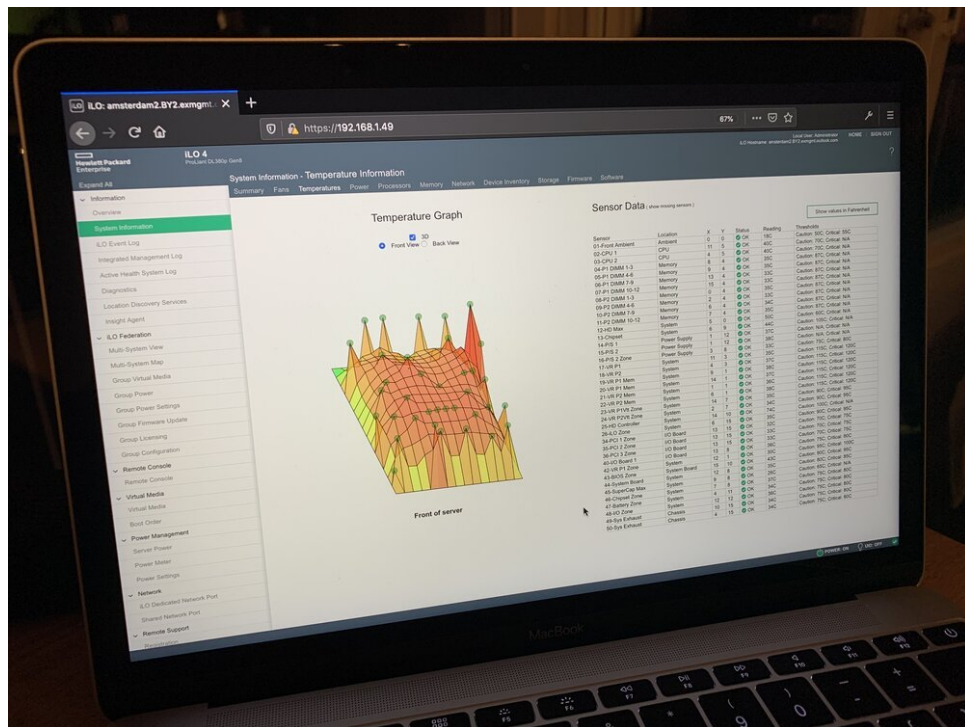
Server (top left) connected to the computers on my desk (lower right) during installation inside my office

Installing the operating system works a little different than your average office computer. The server has no CD / DVD drive to allow installation from physical media. Instead I use a local network to serve the installation files to the server. This is done using my ThinkPad X1 laptop running a webserver. I configured the HP server to connect to this laptop as boot option.

Computers on my desk powering the installation from left to right: ThinkPad X1, iPad Pro and MacBook

MacBook acting as monitor and keyboard of the server

The MacBook is connected to the server through HPE's iLO. It allows it to act as monitor and keyboard. This makes it possible to control the installation from my desk (instead of standing next to the noisy server).

iLO: "sea of sensors" enabling an extensive remote insight in the server

Through iLO it is possible to see all the individual components of the server. It allows engineers to check upon the physical status of the server, see how it's doing.

Bringing it to the datacenter

After the main operating system is installed, the server needs to be placed in the data-center. In Amsterdam there are many different datacenters. My company has multiple servers in different locations, limiting the risk of data loss in case of local fire and disaster. Bringing the server to the datacenter is as simple as making an appointment with the network operation center, signing some contracts and jumping in your car.

Driving to the internet: visiting XS4ALL / KPN in Amsterdam

*Never underestimate the bandwidth of a car with 72TB in its trunk (72TB in my Smart Roadster)*

To allow for an easy installation in the datacenter I preconfigured the network settings to match that of the datacenter. If everything works well, this enables "plug and play" of the server. It did, hooray!

Server connected and online inside the datacenter

Once the server is up and running you can connect to it remotely. This allows you to use it as a real part of "the cloud", tapping into it's computing power from a distance.
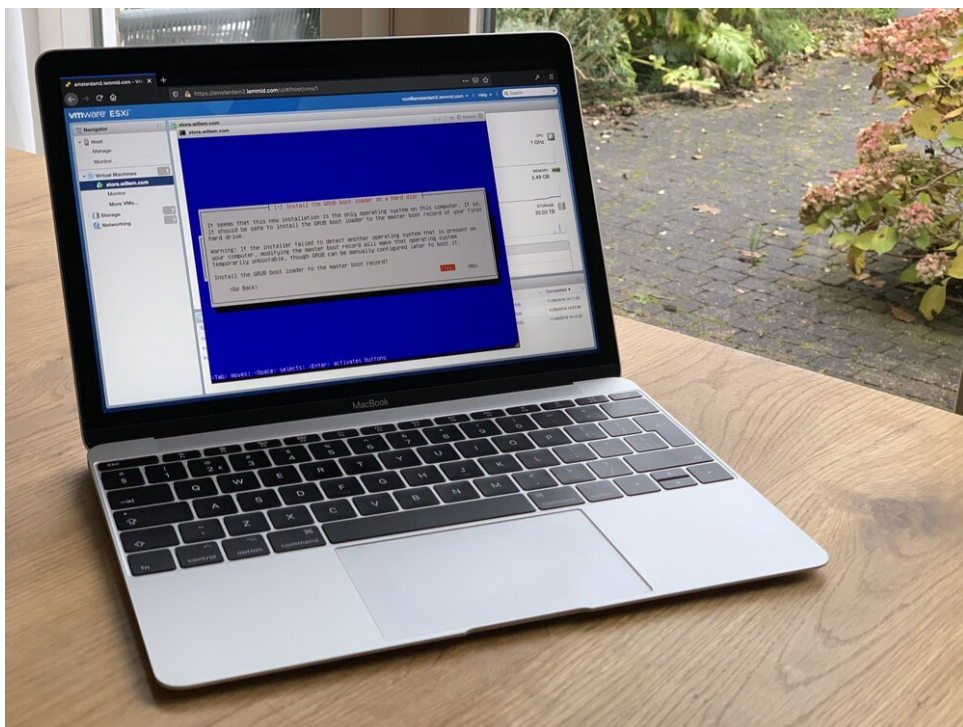
MacBook connected to the server in the datacenter through VMWare ESXi - installing GRUB boot loader on a virtual Debian GNU/Linux server

Conclusion

Powering the cloud are the many server computers inside datacenters. Depending on your needs, you can either rent some capacity from others or install physical machines yourself.

For various reasons I choose to operate my own physical machines, maximising performance and control. Managing servers (both physical and virtual) has become somewhat of a lost art, albeit a critical one for today's cloud powered internet.