



# Nicolás Wolovick



Centro de  
Computación  
de Alto  
Desempeño



Universidad  
Nacional  
de Córdoba

# Machine Learning superserver

- Industria
- Academia

# Nodo Nabucodonosor

- 2 Xeon 2680v2 + 64 GiB RAM + IB QDR + 1 TiB SATA3
- 3 GTX 1080 Ti.



# Founders



I A T E



Se vienen ...



**ASC**  
**ASCENTIO**



# Problemillas ...

## Cryptocurrency miners are renting Boeing 747s to ship graphics cards

By Paul Lilly July 28, 2017

Now that's just 'plane' crazy.

[f](#) [t](#) [g+](#) [e](#) | [r](#) COMMENTS



Have you ever had a moment where you didn't know whether to laugh or cry? That's the situation playing out in the graphics card market because of the cryptocurrency mining boom, a topic we've covered extensively in recent months. But just when we thought there was nothing left to report on the matter, it's come out that some of the most active Ethereum miners are renting Boeing 747 airplanes to ship orders of graphics cards. Yes, seriously.

# Problemillas ...

## Here's why you can't buy a high-end graphics card at Best Buy

"Cryptocurrency can't crash soon enough," one gamer fumes.

TIMOTHY B. LEE - 1/18/2018, 5:35 PM



Matthew Freilich

[Enlarge](#) / Philadelphia miner Matthew Freilich shared this picture of his mining rig, which contains eight Nvidia GTX 1070 graphics cards.



# Problemillas ...



## NVIDIA TITAN V

THE MOST POWERFUL PC GPU EVER CREATED

NVIDIA TITAN V is the most powerful graphics card ever created for the PC, driven by the world's most advanced architecture—NVIDIA Volta. NVIDIA's supercomputing GPU architecture is now here for your PC, and fueling breakthroughs in every industry.

**\$ 2,999.<sup>00</sup>**

**ADD TO CART**

**Free Shipping**

Limit 2 per customer

 **WATCH FULL VIDEO**






# Problemillas ...

Emergent Tech ▶ [Artificial Intelligence](#)

## Nvidia: Using cheap GeForce, Titan GPUs in servers? Haha, nope!

Nice try, but no, you're gonna have to cough up for these expensive data center chips

By [Katyanna Quach](#) 3 Jan 2018 at 03:39

95  [SHARE](#) ▼



Nvidia has banned the use of its GeForce and Titan gaming graphics cards in data centers – forcing organizations to fork out for more expensive gear, like its latest Tesla V100 chips.

The chip-design giant updated its [GeForce and Titan software licensing](#) in the past few days, adding a new clause that reads: "No Datacenter Deployment. The SOFTWARE is not licensed for datacenter deployment, except that blockchain processing in a datacenter is permitted."

The high-end GeForce GTX 1080 Ti graphics cards – aimed at gamers rather than deep-learning data scientists – uses Nv's Pascal architecture, and only costs \$699 (~£514) a pop. Meanwhile, the latest [Tesla V100](#) card that's flogged to data centers costs over \$9,000 (~£6,620).

2.1.2 Linux/FreeBSD Exception. Notwithstanding the foregoing terms of Section 2.1.1, SOFTWARE designed exclusively for use on the Linux or FreeBSD operating systems, or other operating systems derived from the source code to these operating systems, may be copied and redistributed, provided that the binary files thereof are not modified in any way (except for unzipping of compressed files).

### 2.1.3 Limitations.

**No Modification or Reverse Engineering.** Customer may not modify (except as provided in Section 2.1.2), reverse engineer, decompile, or disassemble the SOFTWARE, nor attempt in any other manner to obtain the source code.

**No Separation of Components.** The SOFTWARE is licensed as a single product. Its component parts may not be separated for use on more than one computer, nor otherwise used separately from the other parts.

**No Sublicensing or Distribution.** Customer may not sell, rent, sublicense, distribute or transfer the SOFTWARE; or use the SOFTWARE for public performance or broadcast; or provide commercial hosting services with the SOFTWARE.

**No Datacenter Deployment.** The SOFTWARE is not licensed for datacenter deployment, except that blockchain processing in a datacenter is permitted.

## 3. TERMINATION

This LICENSE will automatically terminate if Customer fails to comply with any of the terms and conditions hereof. In such event, Customer must destroy all copies of the SOFTWARE and all of its component parts.

**Defensive Suspension.** If Customer commences or participates in any legal proceeding against NVIDIA, then NVIDIA may, in its sole discretion, suspend or terminate all license grants and any other rights provided under this LICENSE during the pendency of such legal proceedings.

## 4. OWNERSHIP

All title and intellectual property rights in and to the SOFTWARE (including but not limited to all images, photographs, animations, video, audio, music, text, and other information incorporated into the SOFTWARE), the accompanying materials, and any copies of the SOFTWARE, are owned by NVIDIA or its suppliers.

## 5. APPLICABLE LAW AND BINDING ARBITRATION



**Governing Law.** This LICENSE shall be deemed to have been made in, and shall be construed pursuant to, the laws of the State of Delaware, without regard to or application of its conflict of laws rules or principles. The United Nations Convention on Contracts for the International Sale of Goods is specifically disclaimed.


A pesar de todo, funciona.

```
Fri Feb 23 17:14:39 2018
-----
NVIDIA-SMI 384.111                Driver Version: 384.111
-----
GPU   Name           Persistence-M| Bus-Id           Disp.A | Volatile Uncorr. ECC
Fan  Temp    Perf  Pwr:Usage/Cap|                               Memory-Usage | GPU-Util  Compute M.
-----+-----+-----+-----+-----+-----+-----+-----
  0   GeForce GTX 108...  Off | 00000000:02:00.0 Off |           N/A
    49%   73C    P2  236W / 280W | 1829MiB / 11172MiB |    97%    Default
-----+-----+-----+-----+-----+-----+-----
  1   GeForce GTX 108...  Off | 00000000:03:00.0 Off |           N/A
    23%   63C    P2  233W / 280W | 1829MiB / 11172MiB |    96%    Default
-----+-----+-----+-----+-----+-----+-----
  2   GeForce GTX 108...  Off | 00000000:84:00.0 Off |           N/A
    34%   68C    P2  263W / 280W | 1829MiB / 11172MiB |    96%    Default
-----+-----+-----+-----+-----+-----+-----

Processes:
GPU   PID      Type  Process name                      GPU Memory
-----+-----+-----+-----+-----+-----
  0   13274    C     ./hashcat64.bin                    1819MiB
  1   13274    C     ./hashcat64.bin                    1819MiB
  2   13274    C     ./hashcat64.bin                    1819MiB
-----+-----+-----+-----+-----+-----
```

# Cierre

- 30 TFLOPS single precision.
-  ya estuvo entrenando 
- CTC obtuvo 60K ARS para comprar 1 o 2 Titan V.
- En breve [nabucodonosor.ccad.unc.edu.ar](http://nabucodonosor.ccad.unc.edu.ar)  
¡Pedí tu cuenta!
- Ya estamos pensando en armar el 2<sup>do</sup> y conectarlas por IB.
- Todo esto es obra de Alejandro y Ezequiel




# Preguntas incómodas

- ¿Porque no compran una DGX-1 o una X10DRG-O?
  - ¿Por qué no aquilan una p3.16xlarge en Amazon EC2?
  - ¿Por qué no compran un carrito de minado?
  - ¿Por qué desperdician nodos de Mendieta en esto?
  - ¿Por qué no minan criptomonedas?
- 
- ¿Por qué no nos dejamos de preguntar y entrenamos modelos para darle valor a nuestra ciencia y nuestros productos?



# Problemas resueltos (aka, ¡Entran 4!)



# Problemas resueltos (aka, fan encima de *nasty platters of spinning rust*)

```
anaconda root@localhost /]# hdparm -a /dev/sdb
/dev/sdb:
  readahead = 256 (on)
anaconda root@localhost /]# hdparm -tT /dev/sdb
/dev/sdb:
Timing cached reads: 22856 MB in 2.00 seconds = 11437.97 MB/sec
Timing buffered disk reads: 168 MB in 3.06 seconds = 54.84 MB/sec
anaconda root@localhost /]# hdparm -tT /dev/sda
/dev/sda:
Timing cached reads: 2 MB in 4.50 seconds = 454.68 kB/sec
Timing buffered disk reads: 4 MB in 10.17 seconds = 402.70 kB/sec
anaconda root@localhost /]# hdparm -tT /dev/sda
/dev/sda:
Timing cached reads: 15846 MB in 2.00 seconds = 7928.05 MB/sec
Timing buffered disk reads: 6 MB in 13.61 seconds = 451.55 kB/sec
anaconda root@localhost /]# hdparm -tT /dev/sdb
/dev/sdb:
Timing cached reads: 21528 MB in 2.00 seconds = 10775.88 MB/sec
Timing buffered disk reads: 176 MB in 3.07 seconds = 57.41 MB/sec
anaconda root@localhost /]#
```

```
root@debian:~# cat /[ 611.632946] print_req_error: I/O error, dev sda, sector 2996224
[ 611.634206] Buffer I/O error on device sda1, logical block 374272
[ 611.635309] Buffer I/O error on device sda1, logical block 374273
[ 611.636418] Buffer I/O error on device sda1, logical block 374274
[ 611.637502] Buffer I/O error on device sda1, logical block 374275
[ 611.638570] Buffer I/O error on device sda1, logical block 374276
[ 611.639633] Buffer I/O error on device sda1, logical block 374277
[ 611.640664] Buffer I/O error on device sda1, logical block 374278
[ 611.641665] Buffer I/O error on device sda1, logical block 374279
[ 611.642620] Buffer I/O error on device sda1, logical block 374280
[ 611.643570] Buffer I/O error on device sda1, logical block 374281
[ 611.644668] print_req_error: I/O error, dev sda, sector 3002640
[ 611.645313] print_req_error: I/O error, dev sda, sector 3004416
[ 611.645667] print_req_error: I/O error, dev sda, sector 3006376
[ 611.646275] print_req_error: I/O error, dev sda, sector 3010360
[ 611.646609] print_req_error: I/O error, dev sda, sector 3008512
[ 611.647227] print_req_error: I/O error, dev sda, sector 3012608
[ 611.647542] print_req_error: I/O error, dev sda, sector 3014704
[ 611.648111] print_req_error: I/O error, dev sda, sector 3016704
[ 611.648421] print_req_error: I/O error, dev sda, sector 3020800
```

