

Beyond Desktop Computation

Challenges in Scaling a GPU Infrastructure

Martin Uray¹, Eduard Hirsch¹, Gerold Katzinger², Michael Gadermayr¹

¹ Salzburg University of Applied Sciences, Austria

² NTS Solutions, Salzburg, Austria

International Data Science Conference, October 2021



Industry "Standard" (or Illusion)

Removed slide, due to missing licenses.



SME "Standard" (Figurative Example)



Bob

Data Scientist and Swiss Army Knife



FH Salzburg

Current Setup



Bob
Data Scientist




Bob's Notebook



Computing Server




Problem 1



Solution?



Computing Server




CS 1

CS 2



Problem 2



Problem 2 - Fair Share Scheduling

```
$ nvidia-smi
```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
NVIDIA-SMI		440.33.01		Driver Version: 440.33.01			CUDA Version: 10.2		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
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 RTX 208...	On		00000000:1A:00.0	Off			N/A	
27%	30C	P8	11W / 250W		10651MiB / 11019MiB		71%	Default	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
...									
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
7	GeForce RTX 208...	On		00000000:B2:00.0	Off			N/A	
68%	66C	P2	196W / 250W		10671MiB / 11019MiB		69%	Default	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
Processes:								GPU Memory	
GPU	PID	Type	Process name				Usage		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
0	35622	C	python				10659MiB		
...									
7	38274	C	python				10659MiB		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									



Problem 3



Problem 99



Bob



Problem ∞



Bob




Current Setup - Computing Server

GPU	8x NVIDIA GeForce RTX 2080Ti
CPU	2x Intel Xeon Silver 4114, 2.20 GHz
RAM	230 GB
Storage	230 GB system + 8 TB data



IT Programs @ SUAS



IT Programs @ SUAS

ITS	AISP	BIN	Research	Sum
+100%	15/Year	15/Year	+50%	add. $\approx 250\%$




Requirements

- ▶ Ease of usage
- ▶ Scheduling of Jobs
- ▶ Workload Distribution
- ▶ Permission Management
- ▶ Maintainability and Scalability
- ▶ Network Speed
- ▶ Costs



Cloud?



FH Salzburg


Costs - Comparison On-Premises vs. Cloud

How to evaluate? Logging!


- ▶ Observed four months (Mar.-Jun.)
- ▶ Comparison with current cloud prices
- ▶ naive assumptions



Costs - Comparison On-Premises vs. Cloud



Cloud?




FH Salzburg

On-Premises!

But how?



Computing Infrastructure - Architecture



New Setup

	I	C₁	C₂
GPUs	-	8x NVIDIA GeForce RTX 2080Ti	3x NVIDIA RTX A6000
CPUs	1x AMD EPYC 7302P, 3.00 GHz	2x Intel Xeon Silver 4114, 2.20 GHz	2x AMD EPYC 7452, 2.35GHz
RAM	126 GB	230 GB	512GB
Storage	460 GB system	230 GB system + 8 TB data	240 GB system + 8 TB data



Computing Infrastructure - Software Stack

- ▶ Linux (Ubuntu Server 20.04.3 LTS)
- ▶ shared storage
- ▶ FreeIPA / IPA (User- and Groupmanagement)
- ▶ CUDA
- ▶ SLURM Workloadmanager



SLURM

- ▶ most used Workload Manager¹
- ▶ extensive configuration options
- ▶ scaleable
- ▶ easy to use! Change

```
$ bash start_fancy_deep_learning_model_training.sh
```

to

```
$ sbatch start_fancy_deep_learning_model_training.sh
```

¹<https://top500.org>



Recap: Requirements

- ▶ Ease of usage
- ▶ Scheduling of Jobs
- ▶ Workload Distribution
- ▶ Permission Management
- ▶ Maintainability and Scalability
- ▶ Network Speed
- ▶ Costs



Discussion

- ▶ highly configurable
- ▶ scale-able
- ▶ Higher-Education: strategically smart ?

