

u^b

b

**UNIVERSITÄT
BERN**

u^b

Flexible HPC Billing on UBELIX

Hpc-ch forum, 10.10.2025

Agenda

1. Introduction
2. About UBELIX
3. Timeline of project
4. Development of billing scheme
5. SNSF-approval
6. Stakeholder management
7. Implementation
8. Current state

u^b

Introduction

Team


- 4 people working for UBELIX (330 FTE)
- Operation of hpc-cluster, 2 Storages, Open OnDemand-Plattform and now billing
- Tickets, «full-stack». → Coding-Questions go to DSL
 - 1066 Tickets (YTD)

u^b

Introduction

UBELIX (I)

- **University of Bern Linux Cluster**
- Started in [2001](#), increased development starting 2019
- Currently, funded by “Betriebsmittel” and free* of charge for all users
- Investment-models available



*except storage

u^b

Introduction

UBELIX (II)


- Open to all students, researchers, professors and staff of unibe
- No external usage allowed. Usage requires campus-account.
- currently 1'812 Users with access
- Jobs from ~500 individual users last 3 months
- 5PB of infiniband storage, 75% usage

u^b

Introduction

UBELIX (III)

- 14k CPU-cores --> 238 nodes
 - Type B: 128 cores
 - Type C: 196 cores
- 180 GPUs
 - RTX3090: 5 nodes
 - RTX4090: 9 nodes
 - H100/H200: 8 nodes
- 5PB GPFS storage (infiniband), 6PB nfs-storage



u^b

Billing-scheme

Challenges with being free

- If it's free, nobody cares!
- No motivation for efficiency
- External funding cannot be acquired if there is no price-tag
- Normally, research-resources do not come from administrative-budget!

Timeline

Develop billing scheme

get it approved by SNSF

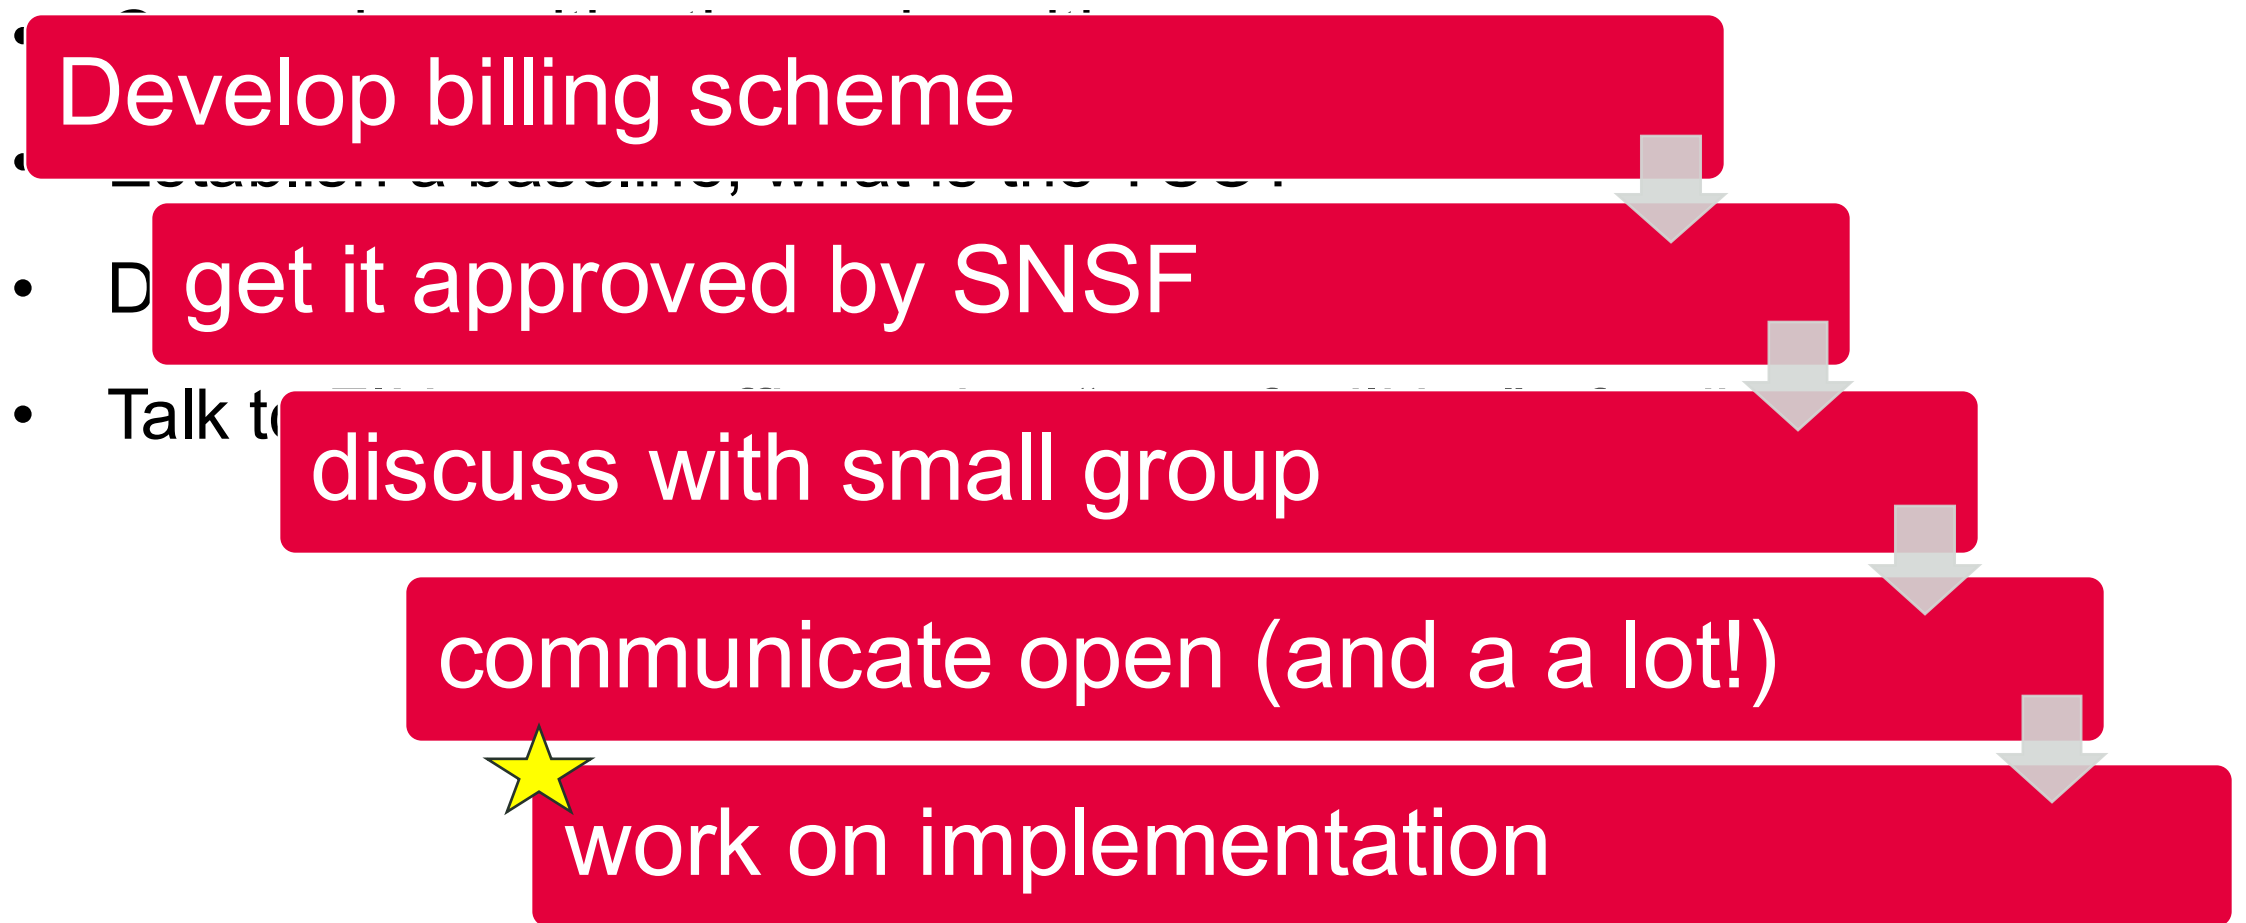
discuss with small group

communicate open (and a a lot!)

 work on implementation

u^b Develop billing scheme

First steps



u^b Develop billing scheme

Result – pay as you go

- $CAPEX / 5 \text{ years} / 8766 \text{ hours}$ → hourly price per server
 - Divided by # resources (GPUs or CPU-cores)
- All supporting-costs (networking, licenses, support, scratch-storage and staff) are fully covered by central IT-services
- Storage is not part of scheme (as of now)
- Invest and free-tiers available

u^b Develop billing scheme

Result – invest

- Contract duration 1-5 years
- Cost: $CAPEX / 5 \text{ years} * \text{contract-duration}$ minus 6%/yr discount
- Hardware-Warranty: 3 years. 2 additional years “on us”
- “Pseudo exclusive access” → Running jobs being preempted if investor submits.
- Max 40% of resources are sold to investors!

u^b Develop billing scheme

Free-tiers

F1: "User" – Pool of free resources

- 1152 TRES minutes (e.g. 8 cores for 24h). Also including GPUs

F2: "project-group"

- Yearly deduction of 1kCHF per project-group

F3: "preemptable"

- All invested resources can be used for free in preemptable-qos. Checkpointing required!

u^b Develop billing scheme

Timeline

- Start billing after SNSF fall-deadline October 2025
- To help projects already ongoing:
 - Free-tier 2: Deduct 3kCHF in the first year (2026), 2kCHF in 2027 and 1kCHF for following years

Approval by SNSF

- Luckily, quite straight forward!
- Fill standardized pricing-sheet
- Wait for approval
- → no issues!

Annex 2

Operational and unit costs for the HPC-Cluster UBELIX

		CPU	GPU	GPU+ (KI Server)		Researc h Storage	Capacity Storage	
Yearly operational costs, CHF								
Direct costs (eligible)								
1.1 Salaries and social charges of operational personnel		1	8	8	1	6	3	
1.2 Consumables		0	0	0	0	0	0	
1.3 Fluid costs		0	0	0	0	0	0	
1.4 Energy (for equipment where this variable cost is significant)		1	5	3	3	7	4	
1.5 Goods of non-enduring value (≤4 years)			8	15		0	1	
1.6 Others		1				8		
Total direct costs	CHF	3	1	12		10	3	
Other direct costs (non-eligible)								
2.1 Salaries and social charges of maintenance personnel			19	19	9	1	4	
2.2 Maintenance contracts (annual preventive maintenance), repair		0	0	0	0	0	0	
2.3 Depreciation		0	0	0	0	0	0	
2.4 Other costs directly linked to the infrastructure		0	0	0	0	0	0	
Total other direct costs	CHF		9	9	9	14	4	
Indirect costs (non-eligible)								
3.1 Salaries and social charges of administrative personnel		3	4	4	4	2	2	
3.2 Salaries and social charges of central services personnel		0	0	0	0	0	0	
3.3 Rent and running expenses of the platform		0	0	0	0	0	0	
3.4 Other indirect and overhead costs		0	0	0	0	0	0	
Total indirect costs	CHF	3	4	4	4	2	2	
Total yearly operational costs (OPEX)	CHF	994,905	1	4	15	5	1	
			2	.2	.6	TB available	2	
Unit operational costs, CHF			0.002	0.10	0.60	Billed Price per TB/yr	58.00	16.00
U.1 User fee 1 (direct costs)	CHF/h		0.009	0.319	1.11		66.55	24.28
U.2 User fee 2 (direct costs + other direct costs)	CHF/h		0.011	0.40	1.26		68.77	25.84
U.3 User fee 3 (direct costs + other direct costs + indirect costs)	CHF/h		0.012	0.44	1.34		72.10	28.18
Personnel of the common research facility								
P.1 Operational	FTE		190%	120%	70%			
P.2 Maintenance			80%	60%	20%			
P.3 Administrative / central Services			60%	30%	30%			
P.4 Others								
Total			330%					

u^b Stakeholder-mangement

Soundingboard

- Early invitation of key-users
- Very good preparation of presentation (more backup-slides than normal slides!)
- Think about questions upfront (and already answer them)
- In the coffee-break: cluster questions and answer

u^b

Communication

Direct

- Directly discuss open points / questions with members of the soundingboard
- Contact “special requirements customers” to discuss their cases
- Create a FAQ on the internal webpage
- Get invited into courses and meetups

u^b

Communication

”Broadcast”

- Newsletters of grants-office, communication, IT-Services
- Information at monthly “BEKO”
- Mailing to all ”Technikverantwortliche”

u^b

Implementation

Overview

- Meeting with Team of EPFL (merci!) for exchange
- Project-ordering-process using IAM-Portal (One Identity)
- Backend: Elasticsearch
- Project-Members are stored in AD-groups
- Metadata and Jobs are stored in Elasticsearch
- Slurm uses WCKey for Jobsubmission

Request Details ✕

? HPC - Order New Project

Test, HPC (TEST) hpctest ^

Organizational Unit *

This field is mandatory.

User-defined part of Project Name *

Choose HPC Project Space Name*

Not set

Credit Number * i

Project Owner *

Project Administrators * Select

Group Members * Select

modify monthly cost ceiling (CHF) *
0

Project Name Validation

Not set

Credit Number Validation*

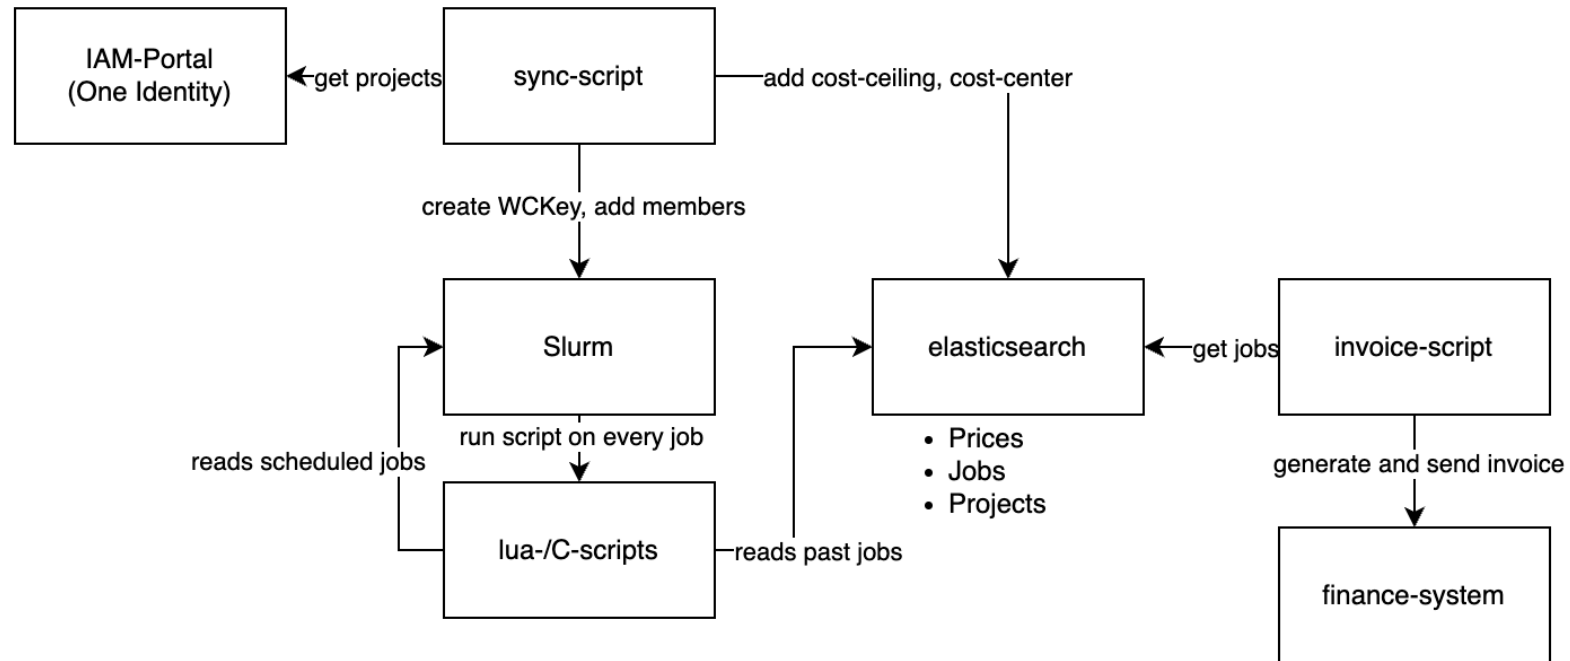
Not set

type 'confirm' to verify input and finalize order *

Skip

Apply

Implementation Workflow



u^b

Implementation

Current state

- Moved go-live-date to 1.1.26
- Issues with load / many small jobs crash plugin / plugin crashes Slurm
- Issues with elasticsearch-plugin: no native auth, not failsafe
- Project-order-workflow ready
- Invoice-Workflow ready

u^b

Implementation

Billing

- Billing is done monthly
- Goal: Automated process, minimal reviewing
- Using internal interface to financing-tool
- No real Invoices, just internal transfer

u^b Future

- Billing
 - Interactive webapp to see current consumption
 - Warning before big job is scheduled
- Infra
 - 2nd cooling-island
 - Scratch on all-flash-system
- Team
 - Hopefully, more staff!

u^b

Questions



Christian Lehmann

Gruppenleiter System Services

Christian.Lehmann@unibe.ch