Leveraging the Jupyter ecosystem for cross-platform deployment of scientific applications

- DESY
- Photon Science
- IDAF
- Maxwell
- Jupyter
- Open Data

Frank Schlünzen – DESY IT 26.02.2024

IDAF/Maxwell (DESY) Frontier (ORNL) Juwels (FZJ)



Accelerators

DESY develops, operates and utilises state-of-the-art accelerator facilities. Scientists from all over the world use these facilities to investigate the structure and function of matter.



B=0

IDAF/Maxwell (DESY)

Photon science

BES

Several of the world's best light sources are located at DESY. Their special X-ray radiation makes atomic structures and reactions in the nanocosmos visible.



IDAF/NAF&Grid

Particle physics

In global cooperations and large teams, DESY scientists investigate the fundamental building blocks and forces of nature.



Astroparticle physics

Astroparticle physics uses various cosmic messengers, such as gamma rays or neutrinos, to understand high-energy processes in the universe.

IDAF/NAF&Grid



Facts and figures



- Executive Summary:
 - Maxwell + Grid + NAF
 - dCache + GPFS + BeeGFS
 - ~60.000 CPU cores, ~320 GPUs
 - ~150 PB data on disk
 - ~2.700 server (compute, storage, management)
 - ~ 0.5 Megawatt

Total Number of compute nodes (CPU + GPU)	950	
Total number of cores with hyperthreading	89256	
Total number of physical cores	44756	
Theoretical CPU peak performance	1548	TFlops
Total RAM	546	тв
Number of GPU nodes	193	
Total number of GPUs	376	
Theoretical GPU peak performance	2615	TFlops
Total peak performance	4163	TFlops
i etal peak periornanee		



HPC'ish

- Infiniband
- Cluster filesystem
- SLURM scheduler
- 4PFlops peak

Less HPC'ish

- Very heterogeneous
- OS migration
- Atypical workload
- Atypical use cases (real time)
- LOTS of partitions
- Non-expert users

Maxwell - Portal V×F=J+∂D/∂

							Ħ	A							Plas			
<u>F</u> ile <u>E</u> dit <u>V</u> iew Hi <u>s</u> tory <u>B</u> ookmarks <u>T</u> o	ols <u>H</u> elp																	- • :
🖻 < 🎲 SSH G what Comput	🕑 sever 🛛 🔽 Zir	mb 🔲 UCla	D 🐞 IDAF	🛞 We×	🚳 comr	🗹 Upda	🥂 Docu 🔒	🗶 dcacl 🛛	💩 File:\ 🛛 🔞 :	Serve Therr	mo comb ir	ne combin e	офь Ргос	Experim	Nova_pa	over 🚳	🞽 P11 b	🗙 Ansj > 🕂 🚿
$\leftarrow \rightarrow C$ \bigcirc \bigcirc \bigcirc	- https://max-por	tal.desy.de/	webavail/										E 🕁		厽	© \$	▽ ± €) 🔃 💷 එ =
M 🕈 🛛 🗶 🋞 Registry 🗹 k 🖗 (G 🕀 🗅 Nvidia+Q) 🗋 Science	🗅 Bikes 🛛	🗅 Haus 🗅	Travel 🗀 D	ashboard 🛛	Personal Ir	nformation .	🗅 DESY	🗅 Fit 🎧 Git	:Hub - dylanre	es/cit 🕀	Share on Co	onfluence				🗅 Other Bookmark
acc-uhh allcpu allcpu-el9 al	lgpu allgpu-el9	cdcs	cfel	cfel-cdi	cfel-cmi	cfel-ux	cssb	cssbgpu	exfel	exfel-th	exfel-theory	exfel-wp72	fspetra	hzg	liv	livgpu	maxcpu	maxcpu-el9
maxgpu maxgpu-el9 mcpu n	npa mpaj	p06	p10	p11	p11x	pcommissi	petra4	petra4-gues	st ponline	ponline_p09	ponline_p11	ponline_p	ps	psgpu	psx	short	topfgpu	uhhxuv
uke upex upex-beam upe	x-high upex-middle	xfel-guest	xfel-op	xfel-sim)													
Timestamp of data (Server): Feb 19 2024 18:08:26 GMT+0100			4						Availab	le					Tota	l nodes 🗌 🖇	Sticky 🔽 To de	etails: click 🧿 hover 🔿
Time since last refresh (local): 00:58										3								
Refresh C Autorefresh:									, i i i i i i i i i i i i i i i i i i i	,								
Colorize selected 🗌 Highlight on hover 🗹																		
																0		
			allgpu						max	gpu						upex		
	Select all Sele	ct none Hidde	en 🗌					(🚯 WebAvail ?)					Tooltij	os 🔽 Animat	ions 🔽 Sort b	y: Partition name 🗸 🗸
				P										Max #	f of Jobs: 80) Max Time	e: 28-00:00:00	Preemption: requeue
			Nodes St	ate Node	elist													
			182 all	nodes max- max-	cdcsg001, max wng[010-067]	x-cfelg[007-011], max-cmsg[0	01-011], max	-cssbg[001-009,	011-014, 018-0)23], max-exflg[009-035], max-	hzgg[003-006	i], max-mpag[i	001-013], ma:	x-p3ag[005-03	35], max-uhhg[001-013],
			2 dra	ained max-	cssbg018, max	x-wng012												
	176		📕 176 all	ocated max- max-	cdcsg001, max wng[010-011, l	x-cfelg007, max 013-067]	-cmsg[001-01	1], max-cssbg	g[001-009, 011-0	14, 019-023], n	nax-exflg[009-0	35], max-hzgg[003-006], max	k-mpag[001-0	13], max-p3aç	g[005-035], m	ax-uhhg[001-0	13],
	170		📕 4 idl	e max-	cfelg[008-011]													
			4 av	ailable max-	cfelg[008-011]													
	3		maxgp	u - partition											Max # of Nodes/ Job: Max # of Jobs	3 Defa	ult Time: 0-01:00	:00 Priority: 50
															Wax # 01 Jubs.	12 M	ax Time: 7-00:00	:00 Preemption: off
			Nodes St	ate N	odelist		[001.013]	wpg[010.06	371									
			1 dra	ained m	ax-wng012	o rij, max-uring	loo i o i ol, ma	whytoro-06	*'1									
			📕 81 all	ocated m	ax-cmsg[001-(011], max-uhhg	[001-013], ma	x-wng[010-01	1, 013-067]									
	81		3 pre	eemptable m	ax-cmsg[003-0	004, 009](partit	ion:allgpu)											
			📕 3 av	ailable m	ax-cmsg[003-(004, 009]												
			upex - j	partition										Max # of No	des/ no limi Job: 200	t Default Time	e: 0-01:00:00	Priority: 500
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0														Max # 01 J	obs. 200	Max Time	e: 14-00:00:00	Preemption: requeue
Node without / with GPU			Nodes St	ate	Nodelist	7. 070 005 5												
■ Idle nodes: 202 21% ■			438 all	nodes I	max-exfl[020-0	71, 076-083, 0	88-099, 101-1	10, 112-150,	152-160, 162-17	9, 181-182, 189	9-360, 369-458,	464-469], max	-exflg[006-00]	7, 009-024, 02	6-027]			
Allocated nodes: 711 75%			1 CO	mpleting i	max-extl434	124 130 152	791											
Other nodes: 37 4%			29 reg	served(idle)	max-exti[009, 1	701(reservation	upex 005719	3 1) max-evfl	[273-280](recerv	ation upex 005	713 2) max-ev	dl[282-290]/reg	ervation	007076) ma	x-exfl[292-29	51(reservation	110ex 900/21	
Iotal number of nodes: 950 Filter	403		403 all	ocated	max-exfl[020-0 009-024, 026-0	071, 076-083, 00 027]	88, 090-099, 1	101-110, 112-	123, 125-138, 14	0-150, 152, 154	4-160, 162-178	, 181-182, 189	-262, 271-272	, 281, 291, 29	6-360, 369-43	33, 435-458, 4	464-469], max-	exflg[006-007,
			0 av	ailable														^

SLURM Rest API

- JWT but without control
- Revoke secret → invalidate all token

Portal serves as token provider

- Revoke indivual token
- Impose JWT lifetime limits
- Allow token delegation



Grafana REST Api

- Only access from individual hosts
- No restrictions host can see all metrics

Portal+WebJobs serve as Grafana API

- Limit access to metrics of interest
- User can see only *own* metrics
- Admin can see it all



Maxwell - WebJobs V×F=J+0D/0-

Username or userID: allusers

Go C

✓ pending(493/493) ✓ running(749/749) ✓ completing(1/r) ✓ completed(254/254) ✓ cancelled(1/r) ✓ failed(1/r) ✓ all(1499/r499) □ none filter: [.... = does not contain Clear

Current jobs for user allusers:

JobID	Job Name	State	Nodes v	Nodelist	Partition	User
5877402	SiAuSi1	running	64	max-exfl[296-359]	extel	
5898092	emc_groel	running	25	max-extl[189,191-193,19	upex	
5903215	emc_groel	pending	25		upex	
5891391	SCIrack	running	10	max-p4-[027-036]	petra4	
5698440	270Ktreeze	running	8	max-p4-[003-010]	allcpu	
5906881	250Ktreeze	running	8	max-terrari[001-008]	allcpu	
5891478	spawner-jupyterhub	running	7	max-extl[141-147]	upex	
5901565	Interactive	running	6	max-exti[039,061-065]	extel	
5884840	13	running	4	max-wn[104-107]	maxcpu	
5898465	spawner-jupyterhub	running	4	max-extl[402-405]	upex	
5900344	Als-dsmc3D	running	4	max-wn[067-070]	allcpu	
5906431	HIPACE	cancelled	4	max-mpag[001,009-011]	mpa	
58/6311	relion40	running	3	max-uhhg[001-002,004]	maxgpu	
5827259	FB-WI	running	2	max-mpag[003,013]	mpa	
5879936	relion40	running	2	max-uhhg[005-006]	maxgpu	
5890417	sext_err	failed	2	max-wng[024,038]	maxgpu	
5890427	sext_err	running	2	max-wng[035,066]	maxgpu	
5890429	sext_err	running	2	max-wng[022,055]	maxgpu	
5898561	spawner-jupyterhub	running	2	max-exfl[043-044]	upex	
5899906	Iomographyh_measured	running	2	max-wn[109-110]	maxcpu	
5097165	dummy_11019431	running	1	max-p3a031	ponline	
5118051	dummy_11019445	running	1	max-p3ag021	ponline	
5194998	dummy_11019459	running	1	max-p3a032	ponline	
5220405	dummy_11019221	running	1	max-p3a033	ponline	
5620106	cssb_cpu	running	1	max-cssb006	cssb	
5734568	comsol-test	running	1	max-p3a030	ps	
5736005	spawner-jupyterhub	running	1	max-wne008	jhub	
5755077	spawner-jupyterhub	running	1	max-wne008	jhub	
5755275	interactive	running	1	max-p3ag019	ps	
5755292	spawner-jupyterhub	running	1	max-wne008	jhub	
5755359	spawner-jupyterhub	running	1	max-wne004	jhub	
5755457	colabfold	running	1	max-cssbg022	cssbgpu	
5755655	spawner-jupyterhub	running	1	max-wne006	jhub	
5757125	spawner-jupyterhub	running	1	max-wne006	jhub	
5760270	spawner-jupyterhub	running	1	max-wne004	jhub	
5770249	spawner-jupyterhub	running	1	max-wne002	jhub	
5770453	spawner-jupyterhub	running	1	max-wne006	jhub	
5772316	spawner-jupyterhub	running	1	max-wne002	jhub	
5772326	spawner-jupyterhub	running	1	max-wne002	jhub	
5772554	spawner-jupyterhub	running	1	max-wne006	jhub	
5778922	spawner-jupyterhub	running	1	max-wne002	jhub	
5788690	spawner-jupyterhub	running	1	max-wne002	jhub	
5789359	spawner-jupyterhub	running	1	max-wne002	jhub	
5798453	spawner-jupyterhub	running	1	max-wne002	jhub	
5801039	spawner-jupyterhub	running	1	max-wne002	jhub	
5801362	spawner-iupyterhub	running	1	max-wne002	ihub	

Job ID: 5891478 Job name: spawner-jupyterhub

 Job state:

 running

 Start:
 3 hours ago
 (Feb 19 2024 14:27:05 GMT+0100 (1708349225))

 Elapsed:
 3 hours, 56 minutes and 16 seconds (14176)

Partition: upex #Nodes: 7 Nodelist: max-exti[141-147] shared node(s): no #GPUs: 0

Add logical cores (HT)

max-exfl141: CPU 2 % memory: 7 %





max-exfl142: CPU 0 % memory: 6 %





max-exfl143: CPU 0 % memory: 6 %

100.%	CPU usage	e (Ø: 0 %)	
100 %			



Maxwell - WebJobs V×F=J+0D/0-

Userna	me or userID:	allusers		Go 🤇			
pendin	g(493/493) 🔽 rur	nning(749/749) 🔽 comple	ting(1/1) 🔽	completed(254/254)	ncelled(1/1) 🗹 failed(1/1)	
all (149	9/1499) []none	filter: ! =	does not cor	ntain	Clear		
Current	t jobs for user a	Ilusers:					
JobID	Job Na	me	State	Nodes v	Nodelist	Partition	User
5877402	SiAuSi1		running	64	max-exfl[296-359]	exfel	
5898092	emc_groel		running	25	max-exfl[189,191-193,19	upex	
5903215	emc_groel		pending	25		upex	
5891391	SCTrack		running	10	max-p4-[027-036]	petra4	
5698440	270Kfreeze		running	8	max-p4-[003-010]	allcpu	
5906881	250Kfreeze		running	8	max-ferrari[001-008]	allcpu	
5891478	spawner-jupyter	hub	running	7	max-exfl[141-147]	upex	
5901565	interactive	-				exfel	
5884840	T3			528	39631 schluenz.pdf	maxcpu	
5898465	spawner-jupyte	(DESY.)		010	_contach2.par	upex	
5900344	Als-dsmc3D					allcpu	
5906431	HIPACE					mpa	
5876311	relion40	Usernam	e: schlue	nz		maxgpu	
5827259	FB-WT	Job ID:	528963	31		mpa	
5879936	relion40	000 10.	020000			maxopu	
5890417	sext err	Job nam	e: spawn	er-jupyter	hub.Maxwell	maxopu	
5890427	sext err	Job state	running:	a		maxopu	
5890429	sext err	<u>.</u>				maxopu	
5898561	spawner-iupvte	Start:	2024-0	01-23 08:5	51:24	upex	
5899906	Tomographyh	Elapsed:	21 min	nutes and	45 seconds	maxcpu	
5097165	dummy 11019	Drintod:	2024.0	1 22 00-1	12-22	ponline	
5118051	dummy 11019	Finited.	2024-0	1-23 09.1	13.33	ponline	
5194998	dummy 11019	Partition:	maxcp	bu		ponline	
5220405	dummy 11019	#Nodes:	1			ponline	
5620106	cssb cpu	intoaco.				cssb	
5734568	comsol-test	Nodelist:	max-w	/n030		ns	
5736005	snawner-junvte	#GPUs:	0	_		ihuh	
5755077	snawner-jupyte	_	-			ihub	
5755275	interactive	Energy:	0.32 k	wn		ns	
5755292	snawner-junkt	CO2:	0.16 k	g /		ibub	
5755359	spawner-jupyter	hub	O rupping		max-wne004	ibub	
5755457	colabfold	100		1	max-cssbq022	cssbanu	
5755655	snawner-junvter	hub		1	max-wne006	ihuh	
5757125	spawner-jupyter	hub		1	max-wne006	ibub	
5760270	spawner-jupyter	hub		1	max-wne004	ihub	
5770249	spawner-jupyter	hub		4	max-wpe002	ibub	
5770453	spawner-jupyter	hub		1	max-wne002	jhub	
5772316	spawner-jupyter	hub		1	max-wne000	ihub	
5772326	spawner-jupyter	hub		1	max-wne002	ihub	
5772554	spawner-jupyter	hub		1	max-wne002	ihub	
5778922	spawner-jupyter	hub		1	max-wneoloo	ihub	
5799600	apawner-jupyter	hub	 running running 	1	max-wheeloz	ihub	
5790250	spawner-jupyter	hub	 running running 	1	max-wneuuz	jhub	
5709452	spawner-jupyter	hub	 running running 	1	max-wneuuz	jhub	
5901020	spawner-jupyter	hub	 running running 	1	max-wneuuz	jhub	
5001039	spawner-jupyter	HUD hub	 running 		max-wneuuz	jnub	
3601362	spawner-jupyter	aun	running	1	max-wheeu2	jnub	

Job ID: 5891478 Job name: spawner-jupyterhub

Job state:

unning
Start:

3 hours ago
(Feb 19 2024 14:27:05 GMT+0100 (1708349225))
Elapsed:

3 hours, 56 minutes and 16 seconds (14176)

Partition: upex #Nodes: 7 Nodelist: max-exfl[141-147] shared node(s): no #GPUs: 0

Add logical cores (HT)

max-exfl141: CPU 2 % memory: 7 %





max-exfl142: CPU 0 % memory: 6 %





max-exfl143: CPU 0 % memory: 6 %

100.%	CPU us	sage (Ø	: 0 %)	
100 %				



Maxwell - JupyterHub- ordinary + 20/2

		WIN						
Μ	axwel	Jup	yter J	ob O	ption	S		
	titions jhub	~		•				Select allowed partitions/reservations
Python Vers	Conda	a Python 3.9	~	•				Select environment
@ Launch mod	lus Jupyte	erLab	~	•				Select junyter look & feel
Number of N	lodes 1 0							
Constraints								
O Environmen								
Chvironmen	t .							
						111.		
						111.		
			Start					Submit batch job
			Start			_#	•	Submit batch job
		Node and	Start d GPU availab	bility				- Submit batch job
Partition	# nodes	Node and # avail	Start d GPU availab # GPUs	pility # P100	# V100	# A100	د	Submit batch job
Partition jhub	# nodes 8	Node and # avail 56/240	Start d GPU availat # GPUs 0	sility <mark># P100</mark> 0	# V100 0	# A100 0		Submit batch job
Partition jhub allcpu	# nodes 8 591	Node and # avail 56/240 145	Start d GPU availab # GPUs 0 0	<mark>sility</mark> # P100 0 0	<mark># V100</mark> 0 0	# A100 0 0	4	- Submit batch job
Partition jhub allcpu allgpu	# nodes 8 591 182	Node and # avail 56/240 145 0	Start d GPU availat # GPUs 0 0 0	bility # P100 0 0	<mark># V100</mark> 0 0	# A100 0 0	4	- Submit batch job
Partition jhub allcpu allgpu cfel	# nodes 8 591 182 19	Node and # avail 56/240 145 0 0	Start d GPU availat # GPUs 0 0 0 0	bility # P100 0 0 0	# V100 0 0 0	# Å100 0 0 0	•	- Submit batch job
Partition jhub allcpu allgpu cfel cssb	# nodes 8 591 182 19 8	Node and # avail 56/240 145 0 0 3	Start d GPU availab maintaine d GPUs availab d GPUs d GPU	sility # P100 0 0 0 0 0 0	# V100 0 0 0 0	# A100 0 0 0 0 0	•	- Submit batch job
Partition jhub allcpu allgpu cfel cssb exfel	# nodes 8 591 182 19 8 438	Node and # avail 56/240 145 0 0 0 3 0	Start d GPU availab d GPUs 0 0 0 0 0 0 0 0 0 0 0 0 0	sility # P100 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0	# A100 0 0 0 0 0 0		- Submit batch job
Partition jhub allcpu allgpu cfel cssb exfel maxcpu	# nodes 8 591 182 19 8 438 119	Node and # avail 56/240 145 0 0 3 0 48	Start d GPU availab d GPUs 0 0 0 0 0 0 0 0 0 0 0 0 0	sility # P100 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0		Submit batch job
Partition jhub allcpu allgpu cfel cssb exfel maxcpu maxgpu	# nodes 8 591 182 19 8 438 438 119 82	Node and # avail 56/240 145 0 0 3 0 48 0	Start d GPU availab d GPUs 0 0 0 0 0 0 0 0 0 0 0 0 0	# P100 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0 0 0 0	•	 Submit batch job Show available resources
Partition jhub allcpu allgpu cfel cssb exfel maxcpu maxgpu mpaj	# nodes 8 591 182 19 8 438 438 119 82 4	Node and # avail 56/240 145 0 0 3 0 0 48 0 0	Start d GPU availab d GPUs 0 0 0 0 0 0 0 0 0 0 0 0 0	bility # P100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0 0 0 0 0 0	•	 Submit batch job Show available resources
Partition jhub allcpu allgpu cfel cssb exfel maxcpu maxgpu mpaj petra4	# nodes 8 591 182 19 8 438 119 82 4 62	Node and # avail 56/240 145 0 0 3 0 0 48 0 0 48 0 0 44	Start d GPU availat # GPUs 0	bility # P100 0 0 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	•	 Submit batch job Show available resources
Partition jhub allcpu allgpu cfel cssb exfel maxcpu maxgpu mpaj petra4 ps	# nodes 8 591 182 19 8 438 119 82 4 62 36	Node and # avail 56/240 145 0 0 3 0 48 0 0 48 0 0 44 0	Start d GPU availat # GPUs 0	bility # P100 0 0 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	•	 Submit batch job Show available resources
Partition jhub allcpu allgpu cfel cssb exfel maxcpu maxgpu mpaj petra4 ps psx	# nodes 8 591 182 19 8 438 119 82 4 62 36 26	Node and # avail 56/240 145 0 0 3 0 48 0 0 48 0 0 44 0 0 0	Start d GPU availab # GPUs 0	bility # P100 0 0 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		 Submit batch job Show available resources
Partition jhub allcpu allgpu cfel cssb exfel maxcpu maxgpu mpaj petra4 ps psx uke	# nodes 8 591 182 19 8 438 119 82 4 62 36 26 16	Node and # avail 56/240 145 0 0 3 0 48 0 0 48 0 0 0 44 0 0 0 16	Start d GPU availab # GPUs 0	bility # P100 0 0 0 0 0 0 0 0 0 0 0 0	# V100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	# A100 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		 Submit batch job Show available resources

Maxwell - JupyterHub- ordinary



- Modified batch-spawner (comet-spawner)
- sudo launches batch-job on behalf of user
- Hard limits on runtime
 - 8h for dedicated partitions
 - 7d for jhub partition
- jhub nodes can run 48 concurrent jupyter jobs
 - CPU load ~0
 - Memory load 50-90%
- Extremely popular (sigh)





Maxwell - JupyterHub- ordinary





11/23 11/26 11/29 12/02 12/05 12/08 12/11 12/14 12/17 12/20 12/23 12/26 12/29 01/01 01/04 01/07 01/10 01/13 01/16 01/19 01/22 01/25 01/28 01/31 02/03 02/06 02/09 02/12 02/18

JupyterHub is quite useful beyond running jupyter notebooks

- Comes with user and session management
- Can proxy all kind of (web)-services

Quite flexible

Quite powerful in combination with frameworks

Started to play \rightarrow serious fun (N.Rahmlow)

see what's feasible and what's not

A faster way to build and share data apps

Streamlit turns data scripts into shareable web apps in minutes. All in pure Python. No front-end experience required.

Deploy on Community Cloud (it's free!)





•••	МуАрр.ру		
1 import stream	nlit as st	• • • My App · Streamlit	
<pre>3 st.write(""" 4 # My first ap 5 Hello *world!</pre>	op ! *		=
6 """)		My first app	
		Hello world!	

- Federated login and MFA with keycloak
- maxapp-spawner and custom templates compose landing page
- sbatch scriplet:
 - jhsingle-native-proxy [...] -- streamlit run maxwell.py [...]
- SLURM takes care of scheduling
- Jhub takes care of users & session & proxying
- Everything runs entirely in user space afterwards
 - Full access to all my data and resources
 - No need to cope with e.g. extended ACLs
- Pure python code
 - Very simple to do





HDF5 Viewer	hdf5 🎽	a month ago	start delete
SciCat Frontend	scicat	a month ago	start delete

SciCat frontend - ~5 lines of python code

Beamtimes beamtimeld startTime endTime instrument Principal Investigator Beamtimes 12345678 2022-08-17T12:00:22 2022-08-18T14:15:22 BL1 Stefan Dietrich My first SciCat Proposal 11012907 2022-06-09T07:00:00 2022-06-13T07:00:00 /petra3/p61 Christopher Beyer Chemical and physical properties ... 1 Instruments 11012614 2022-06-09T07:00:00 2022-06-13T07:00:00 /petra3/p03 Kristian Reck In-situ characterization of nucleati. Datasets 11015352 2022-05-30T07:00:00 2022-06-01T07:00:00 /petra3/p08 Florian Bertram In-House Research, Bertram (P08) 🔎 View 11010808 2020-12-17T07:00:00 2020-12-21T07:00:00 /petra3/p03 Stephan Roth In-House Research, Roth (P03) Helmholtz at P11 - working group ... 11013680 2022-05-25T15:00:00 2022-05-25T23:00:00 /petra3/p11 Michael Kolbe 🦳 Folder 11015427 2022-08-18T23:00:00 2022-08-18T15:00:00 /petra3/p11 Jose Velazquez Garcia Crystallographic analysis of the sol... 11002839 2017-06-08T07:00:00 2017-06-14T07:00:00 /petra3/p05 Joerg Hammel Combined crvo- and soft x-ray atte... 11012776 2022-05-27T07:00:00 2022-05-30T07:00:00 /petra3/p02.1 Minghao Yu In-situ Powder Diffraction: Haloge... INDU-22-L01 - Viva Biotech, long-t... 11015533 2022-08-25T07:00:00 2022-08-25T15:00:00 /petra3/p11 Fan Jiang 11015678 Linda Miallau INDLL 22 K22 - Servier 2022-08-26T07:00:00 2022-08-26T15:00:00 /petra3/p11 11014473 2022-08-26T07:00:00 2022-08-30T07:00:00 /petra3/p08 **Richard Harvey** Investigating the forces driving do.. /petra3/p05 Multimodal Characterization of Li-.. 11014977 2022-08-22T07:00:00 2022-08-26T07:00:00 **Ralf Ziesche**

Show statistics

dffull = pd.DataFrame(response.json())



Maxwell at a glance

Welcome Frank Schluenzen

Various tools require API tokens. To generate tokens we need the password used to encrypt ~/.ssh/tokens.crypt.

- Fetch tokens from encrypted file
 - Portal token to generate SLURM tokens → required to submit batch jobs
 - Tokens for max-jhub and naf-jhub
 - Tokens for FastX graphical sessions
- Not extremely convenient but simple enough and sufficiently secure

Just some basics about "my resources"



iger Your Jobs

🏋 Your Beamtimes

📝 Job Composer

👓 Job View

😕 Apps

? Help

Maxwell at a glance

Welcome Frank Schluenzen

Storage & Processes

Home directory: 21.41GB of 30.0GB used (71.37%) in /home/schluenz

Number of files: 467095

BeeGFS: /beegfs/desy/user/schluenz exists

Documentation: Storage on Maxwell

Open FastX sessions: max-display010

Open firefox sessions: max-display001,max-display010,max-display010

Documentation: FastX on Maxwell

Batch-Jobs & Beamtimes

Batch jobs: to view your batch jobs visit the Your Jobs page

Documentation: Running batch jobs on Maxwell

Beamtime information: to view your beamtimes visit the <u>Your Beamtimes page</u>

Facility status: Petra III, FLASH, Eu.XFEL

Documentation: ASAPS3

Licenses & Miscellaneous

Documentation: License usage

Support: maxwell.service@desy.de

Portal: Maxwell Portal

Pulications: Latest publication acknowledging the Maxwell cluster

Pull license information (slow)

Job performance \rightarrow Same mechanism as JobView using the max-portal API to pull grafana metrics

Maxwell	5941658	xdsapp	тахсри	completed	20/02/2024 13:07	20/02/2024 13:16	00:09:07	
	5941646	AlphaFold	maxgpu	running	20/02/2024 13:06	20/02/2024 13:20	00:14:25	
Your Jobs	5870774	interactive	maxgpu	timeout	17/02/2024 23:19	18/02/2024 00:19	01:00:06	
💥 Your Beamtimes	5870742	test-xds.sh	maxgpu	completed	17/02/2024 23:10	17/02/2024 23:15	00:04:49	
📝 Job Composer	5866945	test-xds.sh	maxgpu	completed	17/02/2024 12:17	17/02/2024 13:11	00:53:31	
•• Job View	Your selected job ' xdsapp ' has	the id 5941658						
🔀 Apps								
	Open In JobView							

Select how you would like the data to be displayed:

show usage and memory of all nodes in two seperate graphs

Show usage and memory in one graph per node

Show usage and memory separatly for each node





Job Composer using templates – keep in mind our experiences users

	×
	Submit or Save the Batch Script
	xdsapp Batch Script
 Your Jobs Your Beamtimes Job Composer Job View 	#//bin/bash #/SBATCHinder=0-08:00:00 #SBATCHinder=1 #SBATCHinder=1 #SBATCHinder_shuenz/dsapp-%j.out #SBATCHornme_ischluenz/dsapp-%j.out #SBATCHornme_ischluenz/dsapp-%j.out #SBATCHpartition=maxcpu #SBATCHpartition=maxcpu #SBATCHport.poh-name=xdsapp export OuTDIR='/home/schluenz/xdsapp' export Spacegroup='92' export unit_cell='57.71 57.71 149.86 90.0 90.0 90.0' export H5_MASTER_FILE='/asap3/petra3/gpfs/p11/2023/data/11017810/raw/autoproctest/rotational_002/autoproctest_rotational_002_master.h5' export USER=schluenz source /etc/profile.d/modules.sh module load maxwell ccp4/8.0 xdsapp
	xdsit.pyspacegroup="\$spacegroup \$unit_cell" -i \$H5_MASTER_FILEdir \$OUTDIR/xdsappdelphi=20reint=3nice=0
	Batch File
	/home/schluenz/batch_xdsapp.sh
	submit

View job output - plots generated on the fly



View log for /home/schluenz/xdsapp/xdsapp/CORRECT.LP



3D view of e.g. PDB files generated from AlphaFold templated batch jobs



3rd party applications launched from within the maxwell batch-job running the jhub proxied streamlit app The NAF-button launches a jupyter server via naf-jhub, which initiates a htcondor job on the NAF



Maxwell Mot-a-binder

BinderHub

allows you to BUILD and REGISTER a Docker image from a Git repository

- connect with JupyterHub
- create a public IP address
- interact with the code and environment within a live JupyterHub instance

Without any authentication



Maxwell Not-a-binder

BinderHub

allows you to BUILD and REGISTER a Docker image from a Git repository

- connect with JupyterHub
- create a public IP address
- interact with the code and environment within a live JupyterHub instance
- Without any authentication

Not exactly what we need/want on the Maxwell cluster

... but a service to demonstrate public workflows/notebooks would still be nice Implemented a streamlit-app as a systemd service Runs in a container under a service account behind a proxy

- Quite secure setup
- Uses the "solaris" partition(s)



Open Data Demonstrator

elect a dataset and run a predefined analysis workflow	
Open Data	⊗ ~
SciCat	⊗ ~
spain	⊗ ~

Dataset Information

Experiment: spain	Scientific Metadata
Source: https://public-data.desy.de/	Battery-Level: 67%
Information: Pictures from Benasque	Lens-Info: 24-240mm f/3.5-6.3
Owner: Fuhrmann	Lens-Model: FE 24-240mm F3.5-6.3 OSS
Owner email: patrick.fuhrmann@desy.de	Field of View: 54.4 deg

there is currently no action defined for this dataset

IMPRESSUM - DATA PRIVACY POLICY

Pulling dataset from scicat instance (https://public-data.desy.de)

👻 Notebook launcher 🛛 🗙	+		
\leftarrow \rightarrow C \textcircled{a}	O A = https://max-data.desy.de	90% ☆ Q Search	◙

Open Data Demonstrator

FLASH	8 ~
FL21	⊗ ~
THz Streaking demo	⊗ ~

Dataset Information

Experiment: THz streaking Beamline: FL21 Instrument: THz streaking Sample: Neon Information: The data was taken as part of the THz streaking photon diagnostic shifts. it contains GMD data and Time of Flight data (recorded with the FLASH GHz ADCs) as well as beam arrival time (BAM) data. We show how to read and interpolate the GMD data, look at the TOF data and sort a delay scan with the BAM	Suggested reading R. Ivanov, I. Bermúdez Macias, J. Liu, G. Brenner, J. Roensch-Schulenburg, G. Kurdi, U. Fruehling, K. Wenig, S Walther, A Dimitriou, M Drescher, I P Sazhina, A K Kazansky, N M Kabachnik and S. Düsterer. <i>Single-shot temporal characterization of XUV pulses with duration from</i> ~10 fs to ~350 fs at FLASHJ. Phys. B: At. Mol. Opt. Phys. (2020). <u>https://doi.org/10.1088/1361-6455/ab9c38</u> Dennis Mayer, Fabiano Lever and Markus Gühr, <i>Data analysis procedures for time-resolved x-ray photoelectron spectroscopy at a SASE free- electron-laser</i> , Phys. B: At. Mol. Opt. Phys. 55, 054002 (2022); <u>https://doi.org/10.1088/1361-6455/ac3e91</u>

odes available: 2

Launch the jupyter notebook

Pulling datasets and actions from json (FLASH data)

Maxwell Mot-a-binder



- Automatically spawn notebook
- No authentication needed
- Runs a container as jovyan user
- Uses JupyterHub REST API
- Separate home per session

Maxwell Mot-a-binder - J+00/0



- Docker running headless VNC with pre-installed CrystFEL application
- Uses tollerort.desy.de (harbor container registry) to pull the image
- Bind-mounting /opendata
- Currently only in DESY Network (opening port 443 and websocket proxying still need to be done)

Credits = 0

 $\nabla \times | = 1$

• Axel Wichmann: max-portal, webavail partition monitoring, API

• Neele Rahmlow: WebJobs and most of the streamlit apps