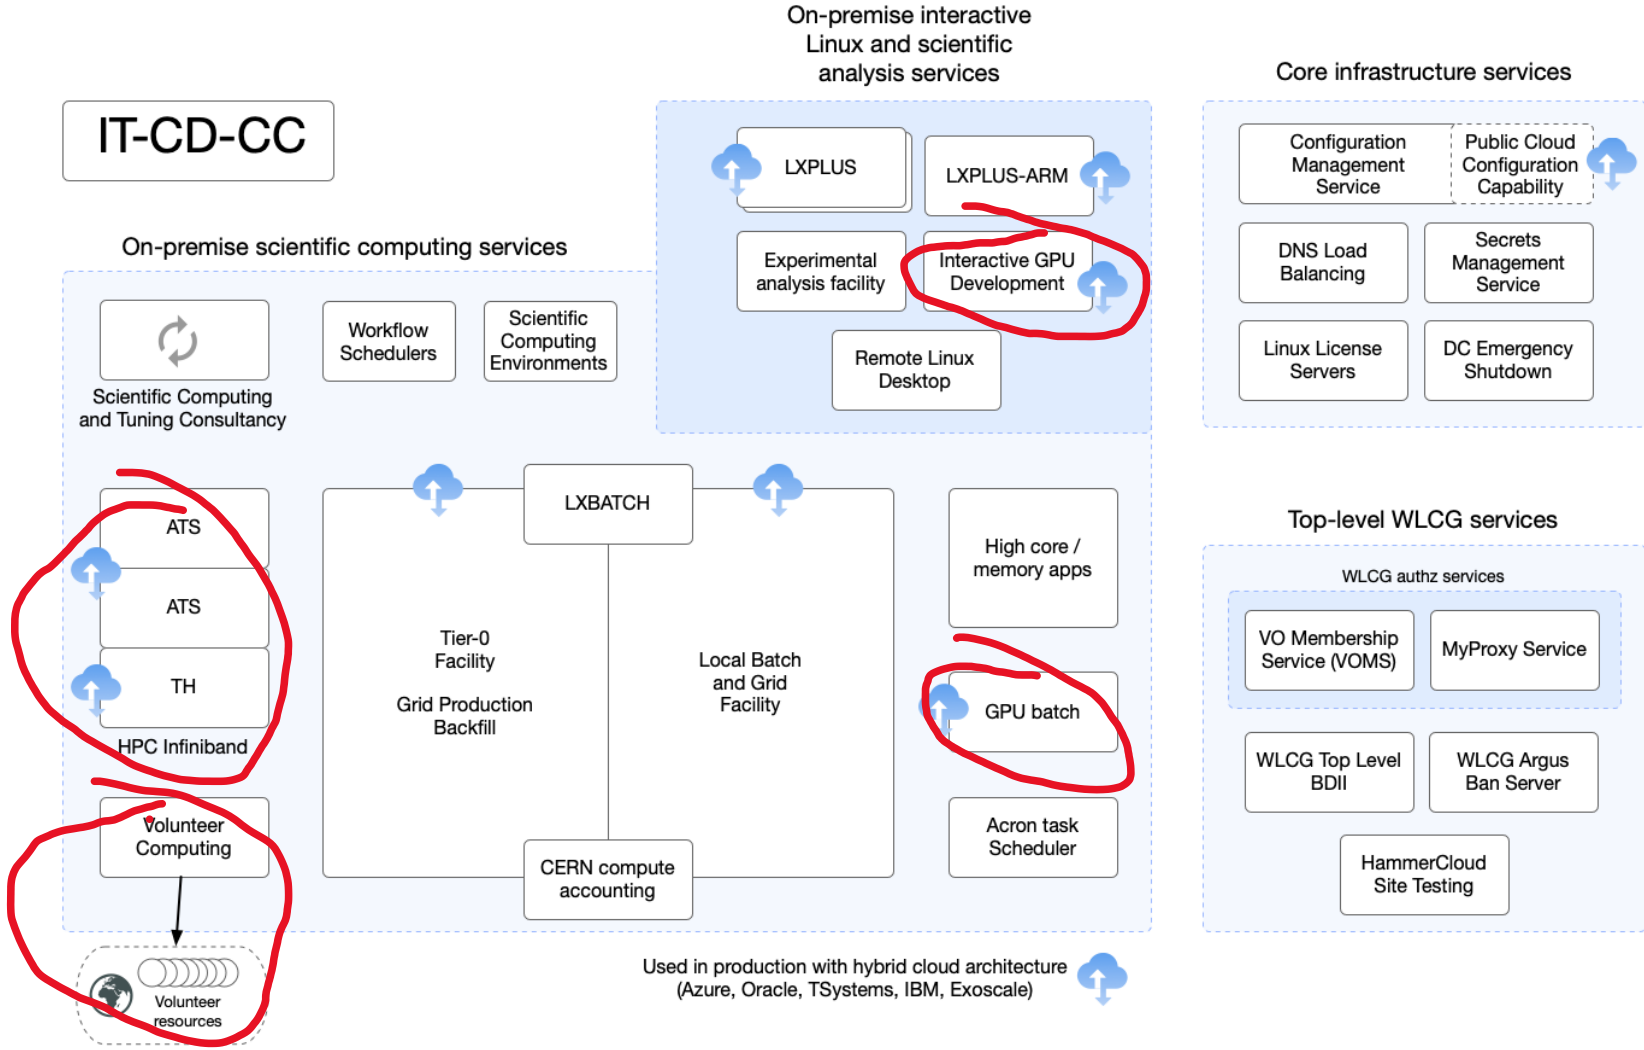


# Compute Services

**Laurence Field**

**CERN IT ML Infrastructure Workshop 10/03/2023**

# Compute and Config



# LxPlus [Linux Public Login User Service]

- **Interactive linux login service for CERN**
- **What is it used for?**
  - Everything and Anything: general purpose computing facility
  - Batch (remote) submit node
  - Users whose primary desktop/laptop is mac or windows do their physics on lxplus
    - Examine subsets of data, prepare jobs, development, LaTeX
  - Remote desktop for graphical apps (vnc, fastx etc)
  - ansible control of their service, Jenkins CI, etc etc
  - Tunneling, email (mutt, alpine)
- **Reference build: people often ask for "their own" lxplus**
  - "contract" lxplus == batch worker node

# LxPlus Current Status

Number of LxPlus 7 Nodes	Number of LxPlus 8 Nodes	Number of LxPlus 9 Nodes
106	18	9




- **Ixplus7 (CERN CentOS7)**
- **Ixplus8 (CentOS Stream 8)**
- **Ixplus9 (CentOS Stream 9)**
- **ixplus.cern.ch alias -> Ixplus7**
- **Ixplus-gpu**
  - **5 Nvidia T4 GPUs**
- **Ixplus node =~ Ixbatch node**
  
- **Active Users**
  - **~1500(day)**
  - **~1000(night)**

# LxBatch

- **Current batch system HTCondor from CHTC Wisconsin**
- **Used for both Grid and "Local" submission**
  - Grid means submitted to a "Compute Element" (CE) which more or less means WLCG
  - "Local" means any user submitting at CERN, authenticated with kerberos
- **High Throughput Computing**
  - "Embarrassingly parallel" or "pleasantly parallel" as HTC would like it to be termed
  - Primary platform for a batch process that can fit on one computer
- **Non-homogeneous resource types**
  - BigMem and BigMCore
  - GPUs
    - 55 A100
    - 23 V100
    - 2-20 T4



# Usage (shared)

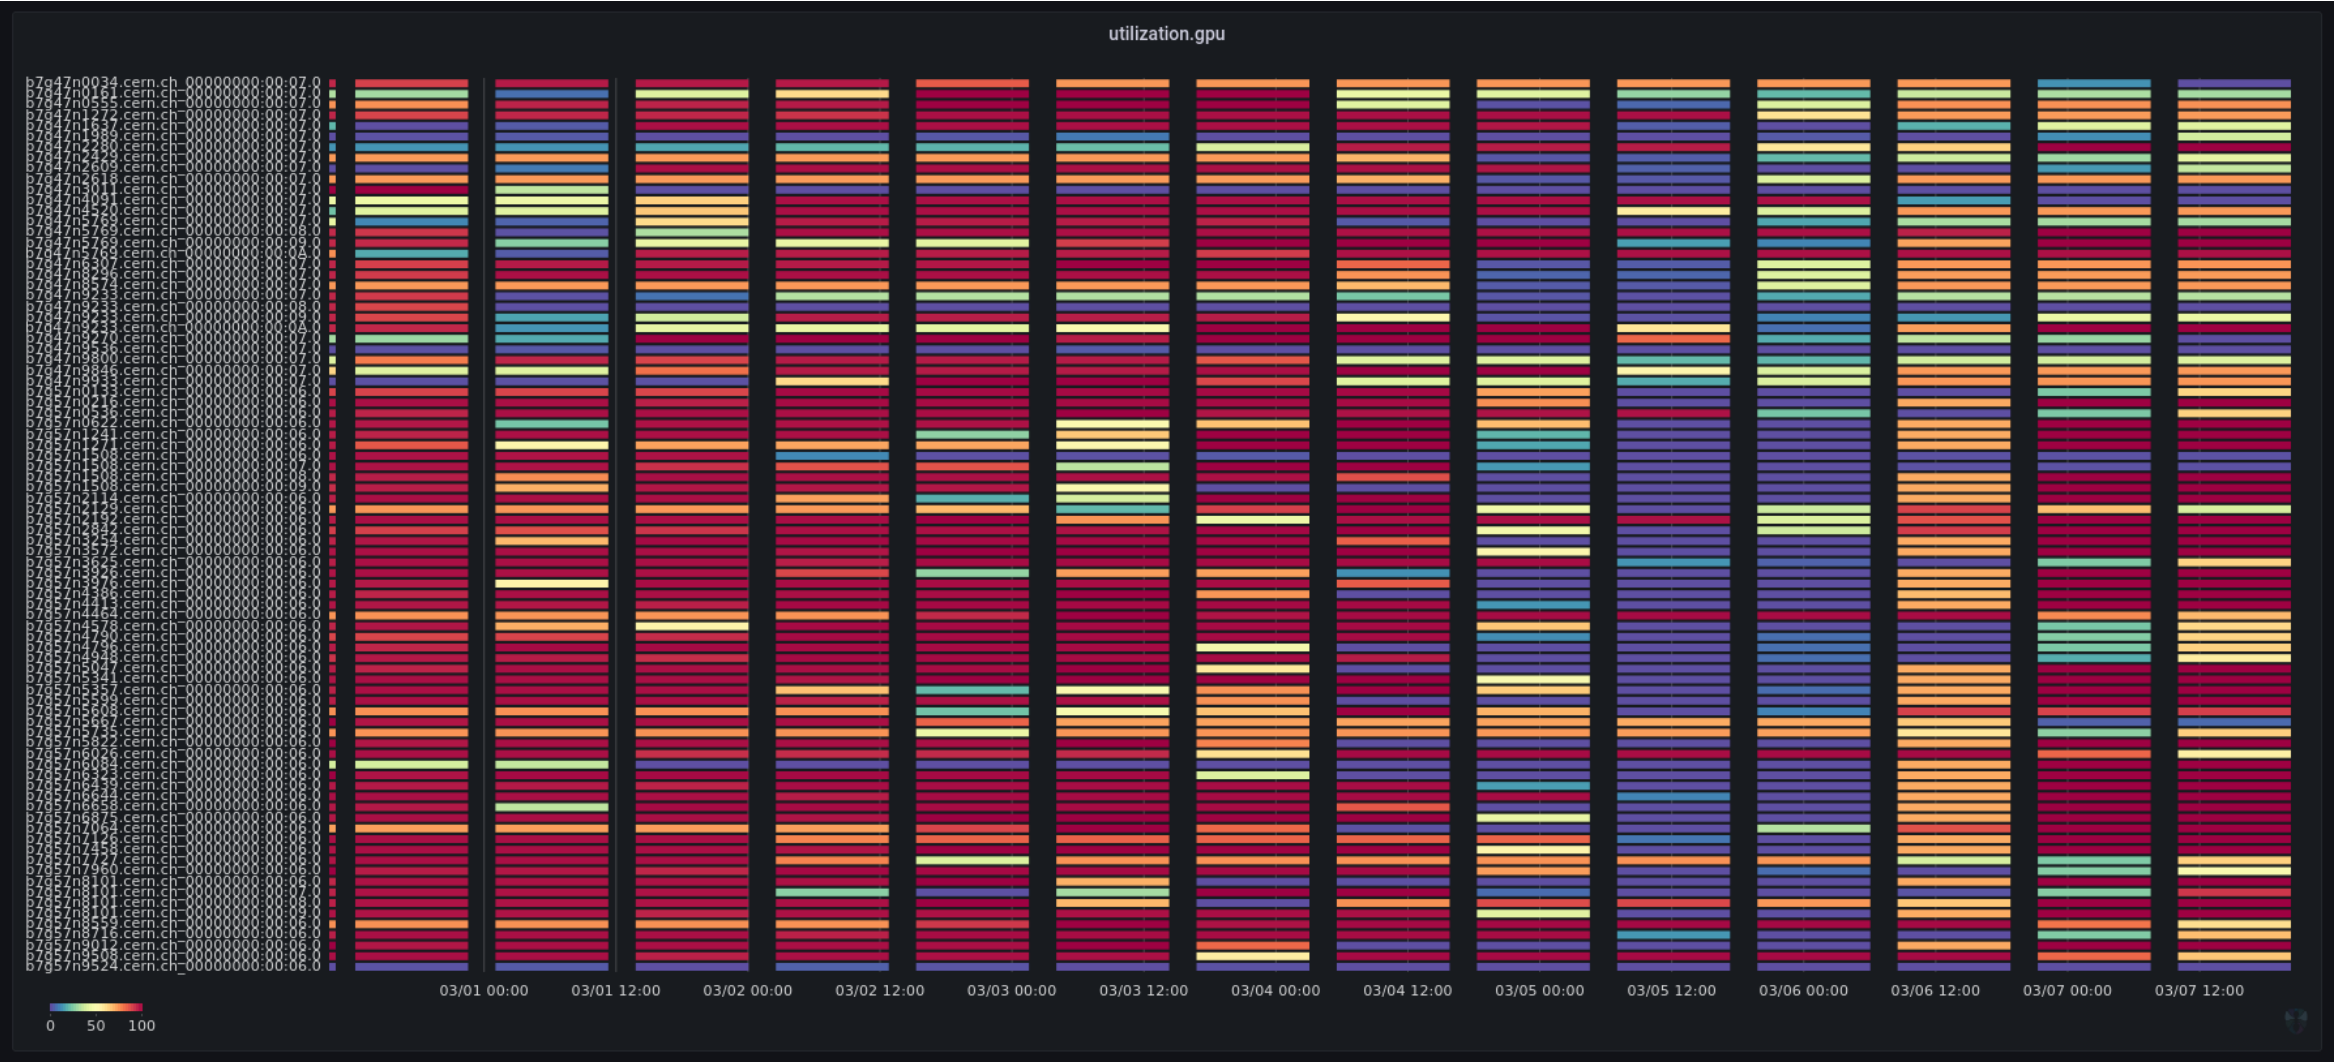


# Usage (dedicated)






# Usage GPUs



# GPU Users

GPU Usage 2022



# High Performance Computing (HPC)

~8000 cores for HPC  
~300 000 cores for batch

- Use cases that do not fit the standard batch HTC model
  - Parallel MPI applications
  - 32-2304+ cores for a single job
- Batch HPC facility using Slurm scheduler to run jobs on 4 Infiniband clusters, on different Slurm partitions:
  - 2x72 nodes with 2 x Xeon(R) CPU E5-2630/20 cores (partitions “inf-short” and “inf-long”)
  - 72 nodes with 2x AMD EPYC 7302 32 cores (partition “photon”)
  - 72 nodes 2 x Xeon(R) CPU E5-2630/16 cores (“qcd” partition)
- Applications that scale well with parallelisation
  - MPI application performance requires fast interconnects with low latency between nodes in a cluster
  - Stability of OS and environment critical
  - Applications typically require fast access to a shared filesystem

# User Community and Applications

- Login to submit node: “hpc-batch.cern.ch”
  - Users' home and scratch directories on /hpcscratch file system (CephFS)
  - AFS and EOS available, similar to Ixplus
  - Applications on AFS or CVMS, (also local or EOS...)
  - EOS for data copy and project storage

## BE

- Plasma simulations for Linac 4
- Beam simulations for LHC, CLIC, FCC...
- Xtrack, PyOrbit etc

## TH

- Lattice QCD simulations

## HSE

- Safety/fire simulations (FDS, OpenFOAM)

## SY

- Gdfdl (field calculations for RF cavities)
- Field calculations (CST...)

## TE


- Picmc
- Engineering (Ansys and Comsol)

## EN

- CFD (Ansys-Fluent, OpenFOAM)
- Structural analysis (Ansys, LS-Dyna...)

# HPC Usage

HPC Usage 2022



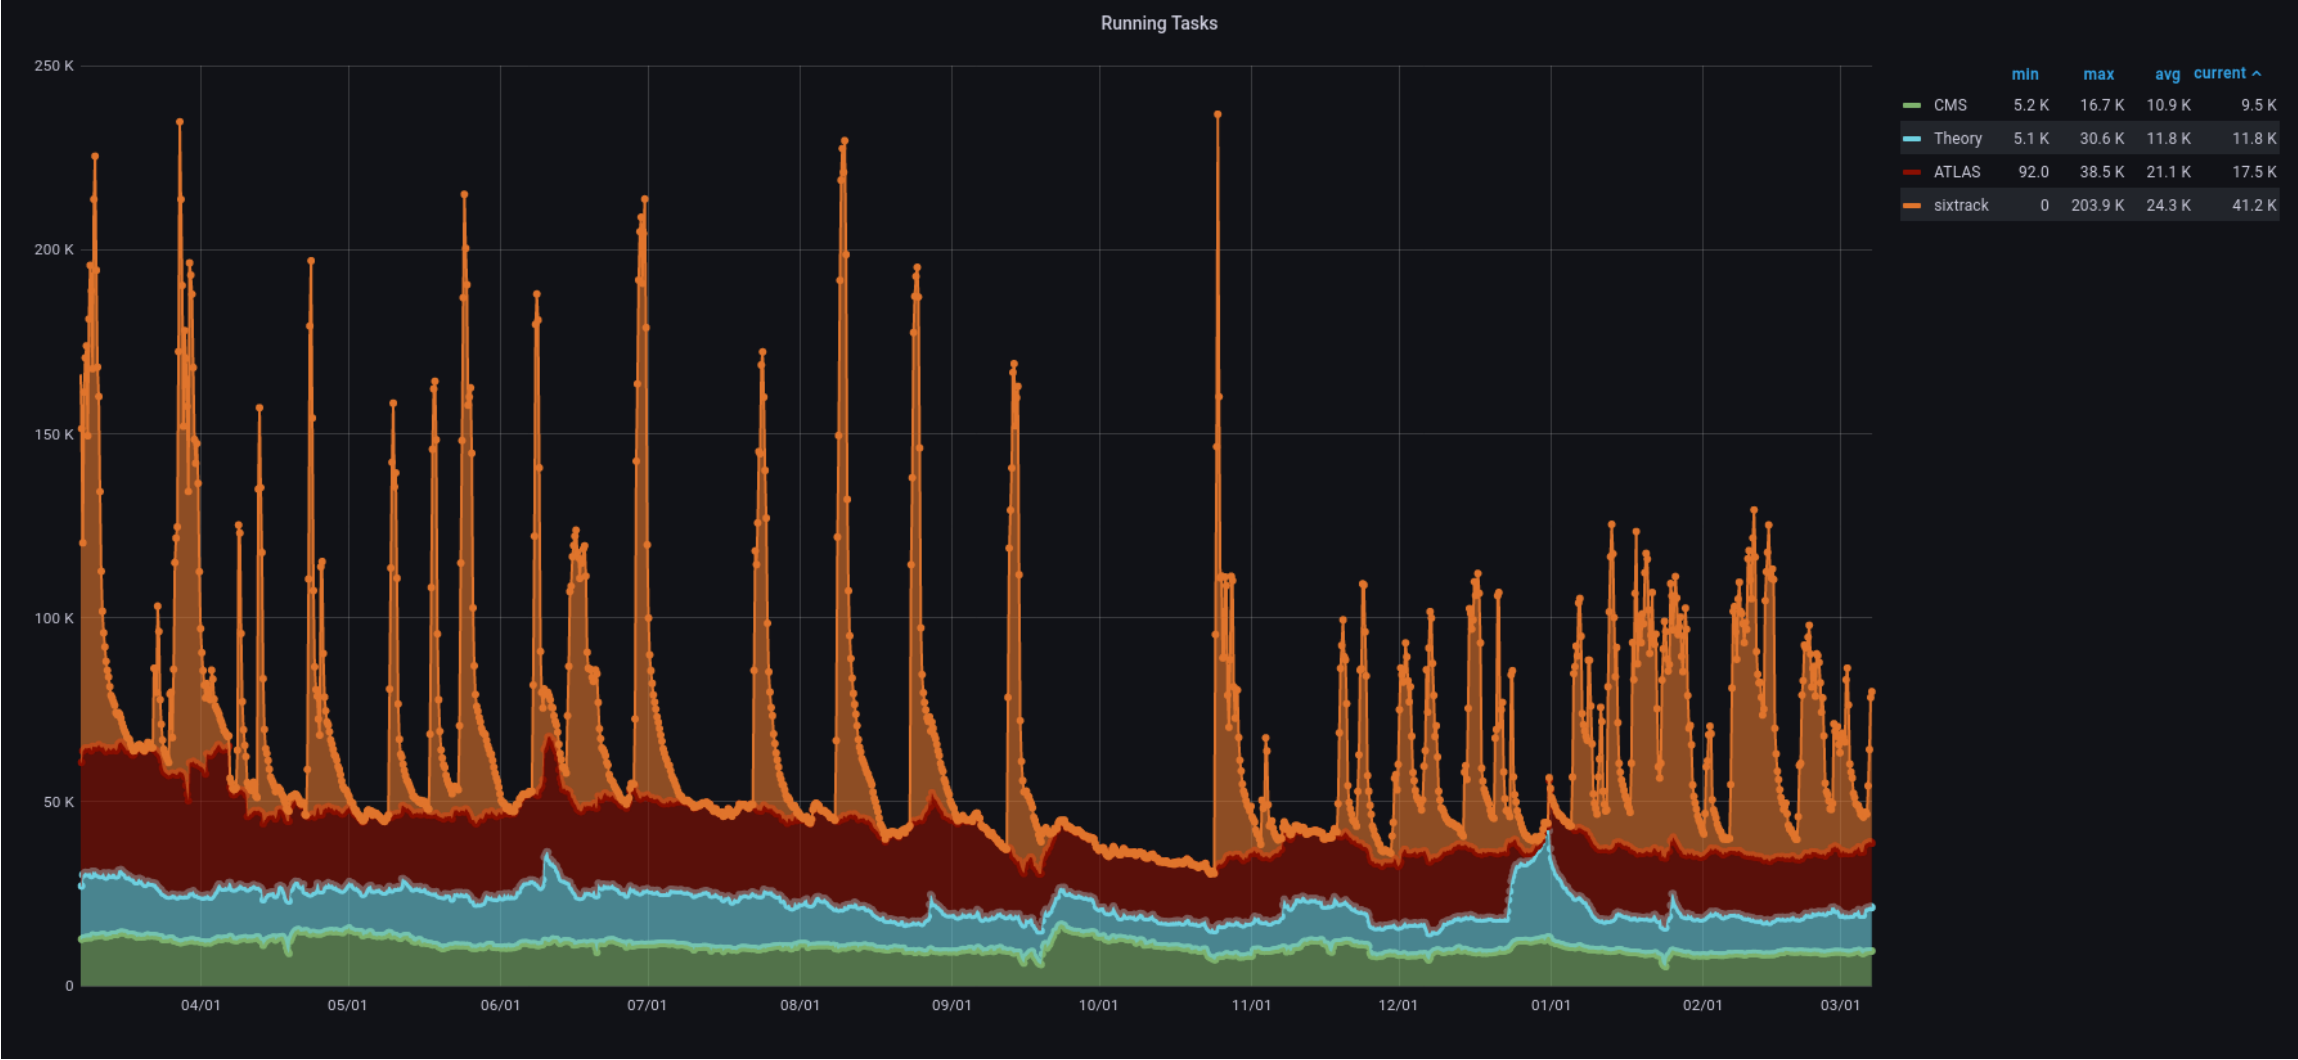
# Overview and history



- **LHC@home**


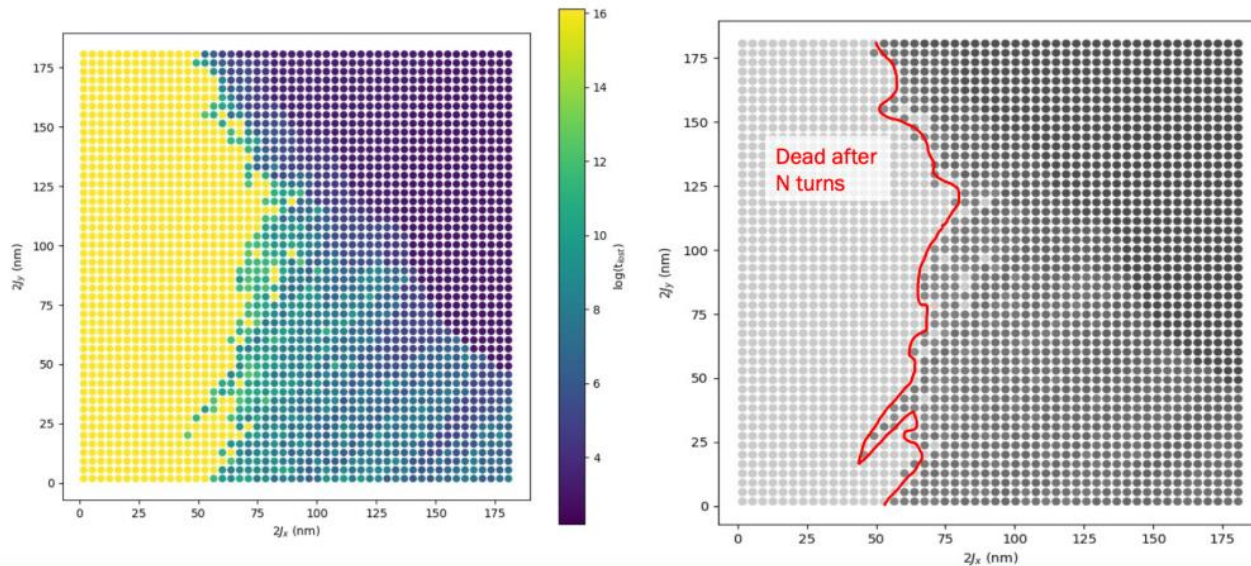
- Started in 2004 for the 50th anniversary of CERN. Running the Sixtrack application for beam simulations of the stability of proton orbits in the LHC accelerator
- From 2011 also physics event simulations (Theory) and later also LHC experiment applications running in a virtual machine (CernVM) on volunteers' computers

# Usage and Opportunity



# Xtrack on BOINC: Example study

- Currently, most people working on LHC simulations already moved to Xtrack (explaining the drop in SixTrack jobs submitted to BOINC)
- The tailoring of SixTrack BOINC to one type of study limited the amount users
- Very large-volume study in the pipeline: surrogate ML model for LHC and FCC in collaboration with EPFL, SDSC (under the CHART project)



Study supported by 



