Prospects for LHC Computing

JFPPL 2016 @ CCIN2P3

Renaud VErnet

OUtline

- Status
- Needs & possible solutions
- Summary

Status

- Higgs boson !
  - And then ?

- « New » Physics

- Why Run II, III, IV and after ?
  - More data ⇨ more sensitivity

Program

	Integrated Lumi (fb
Run 1	25
Run 2	100
Run 3	300
HL-LHC	+300 per year

x 4

x 3

Programme

	Integrated Lumi (fb	Pileup for GPDs
Run 1	25	25
Run 2	100	40
Run 3	300	60
HL-LHC	+300 per year	140

Pile-up

Without

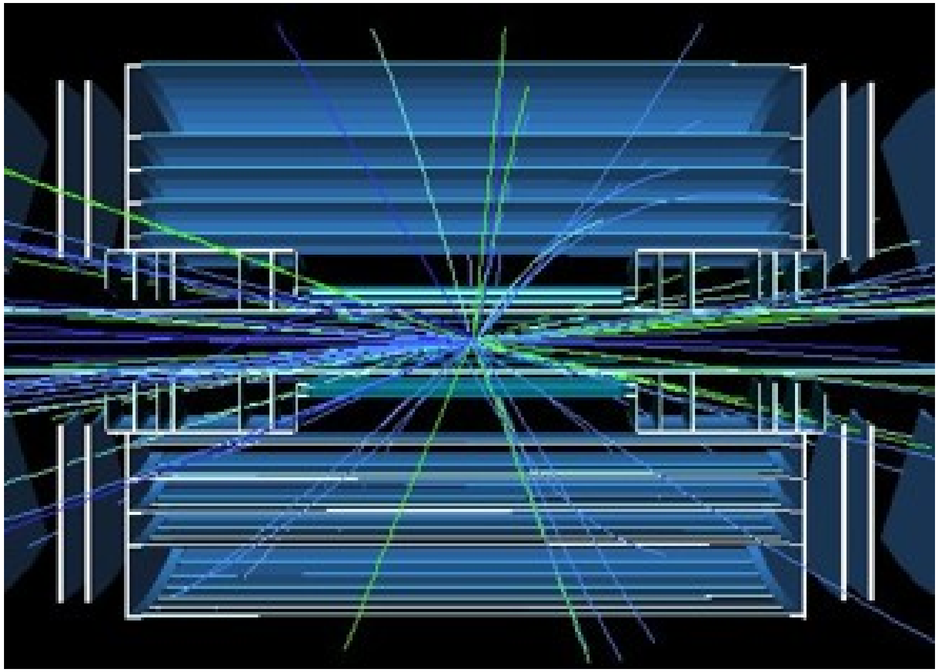

With

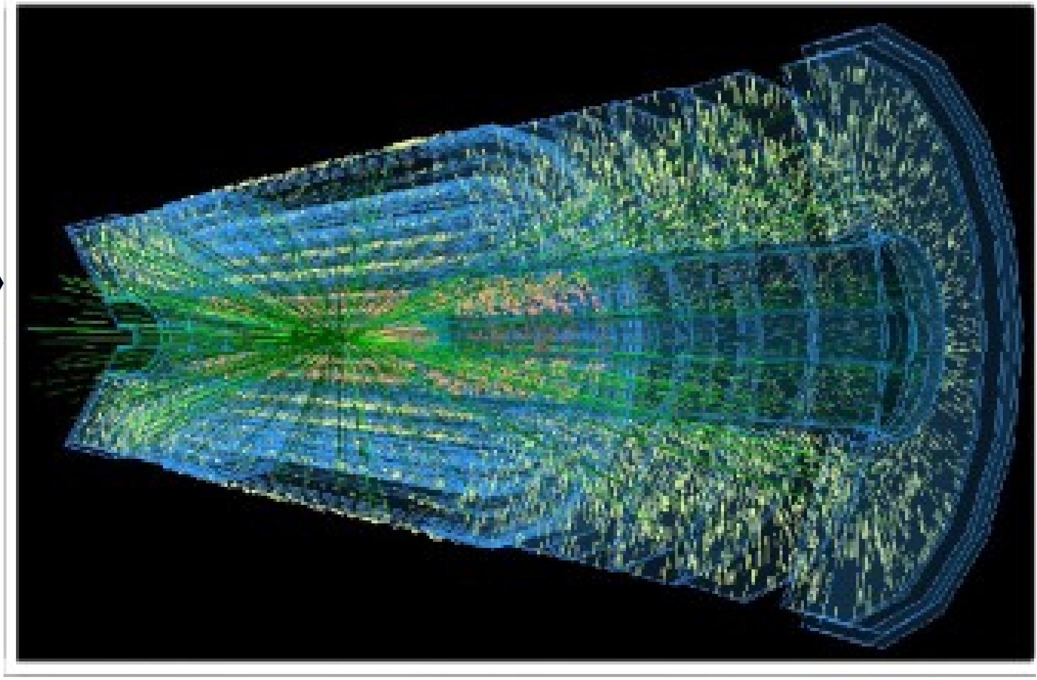

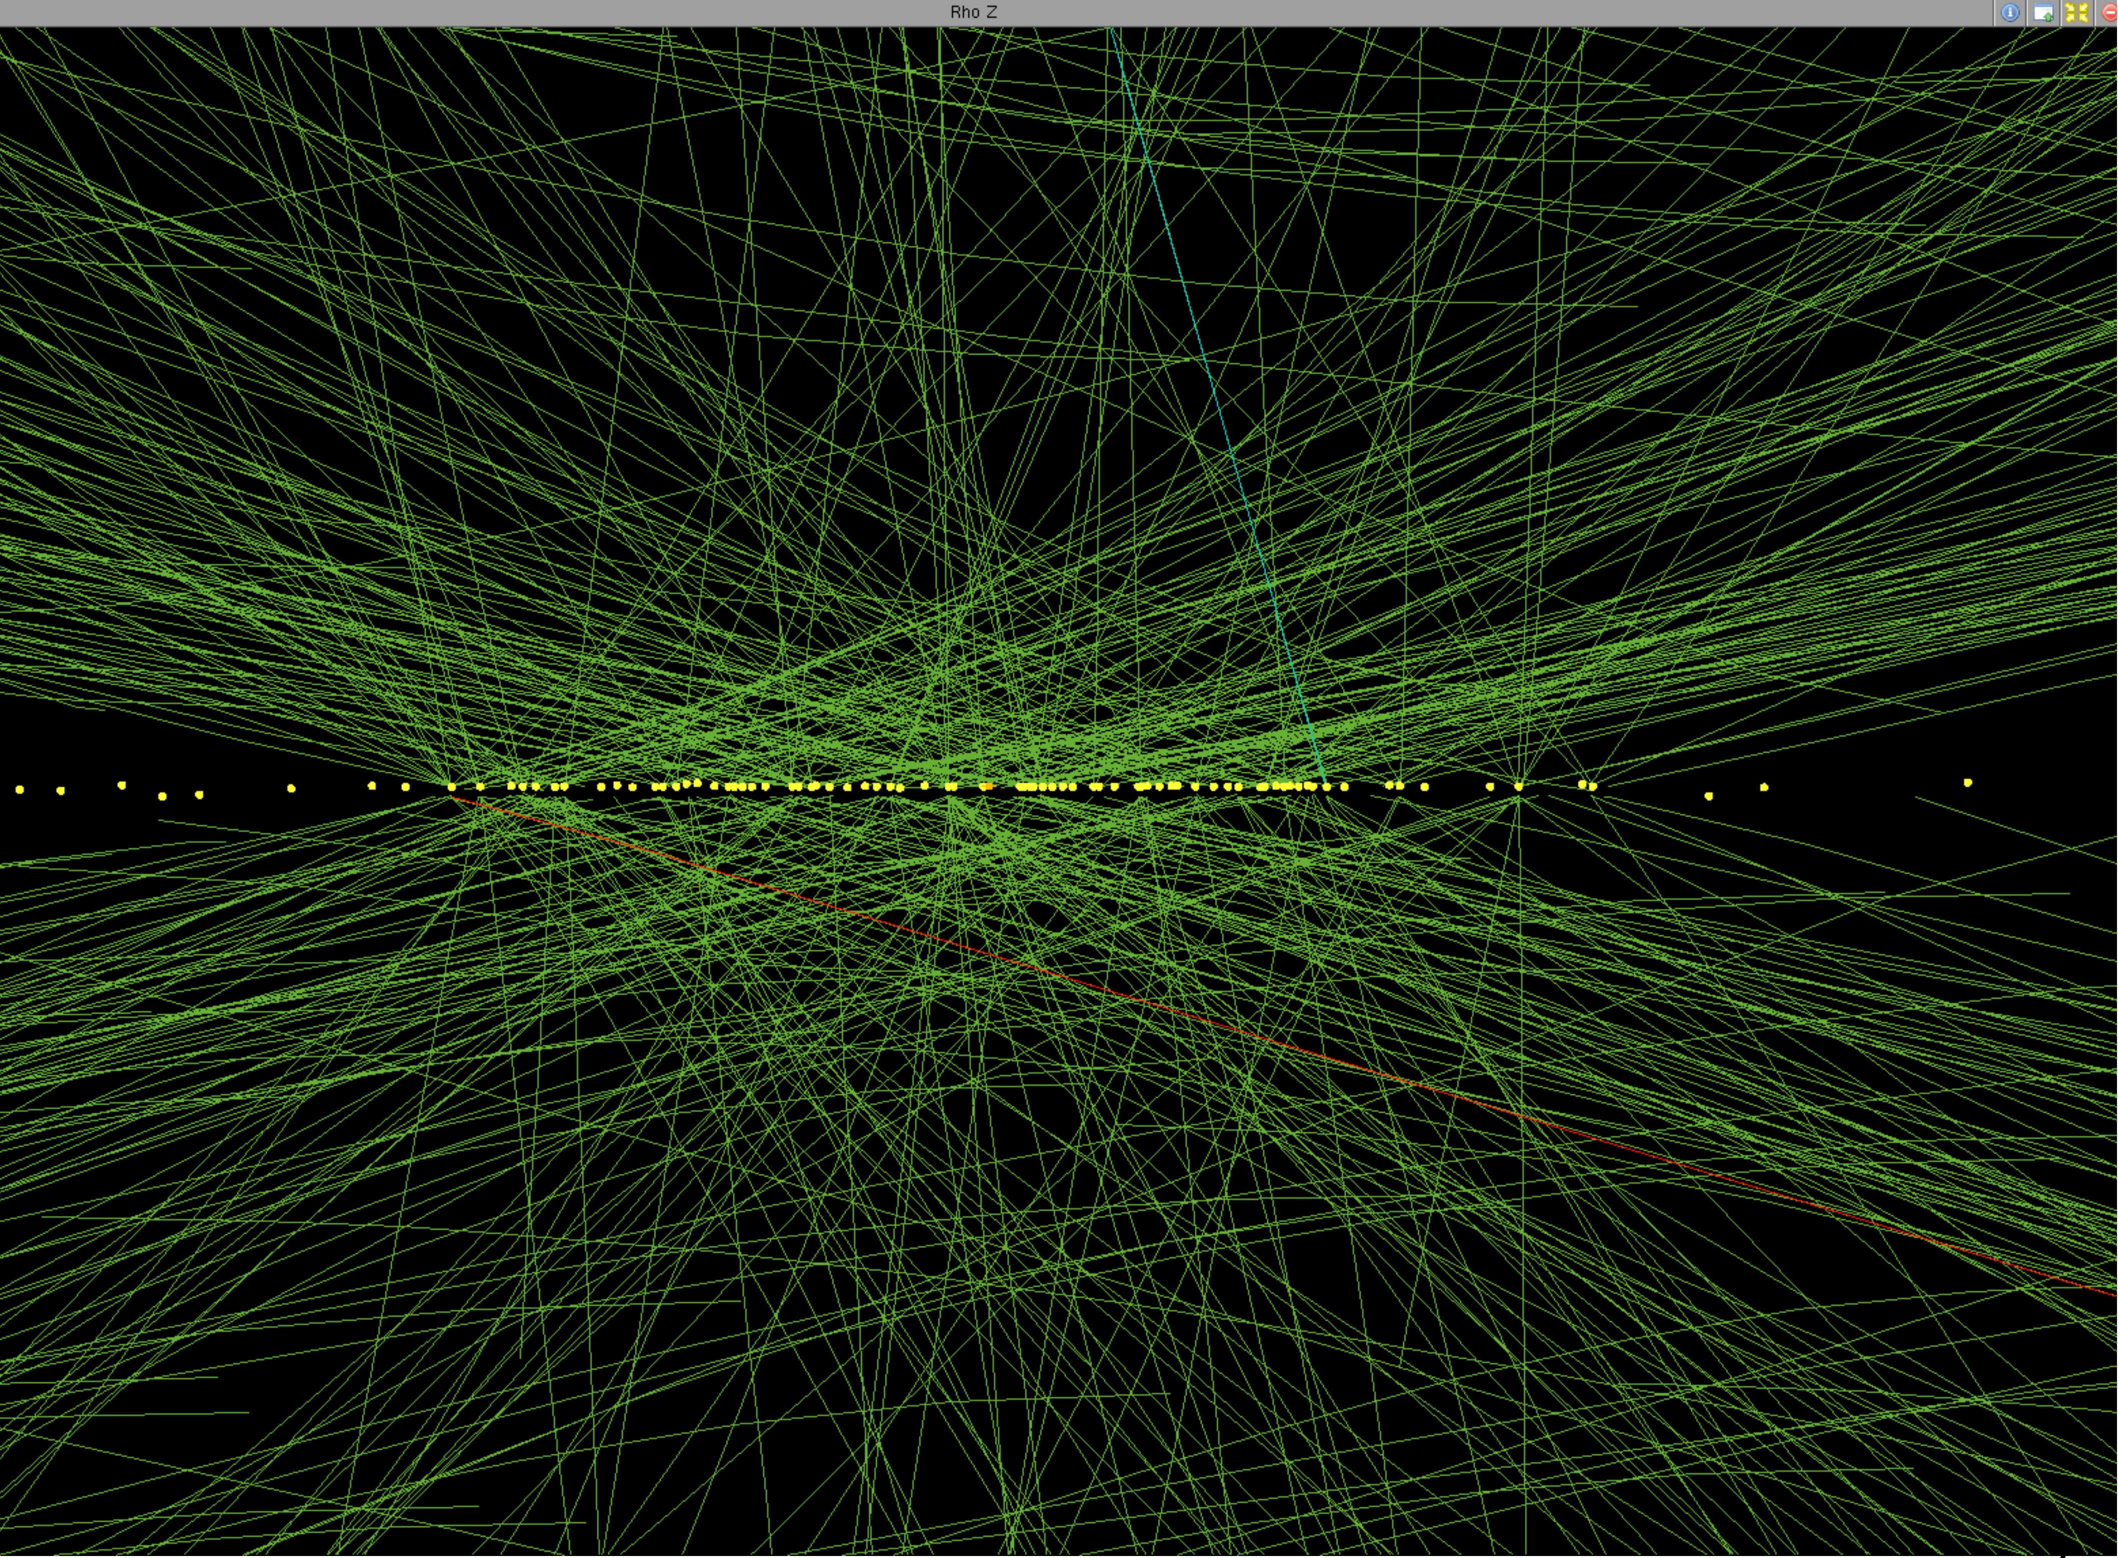
Rho Z

At the end of Run III

- Data x 12
- Pile-up x 2

Big needs

CPU
Storage
Network

}

Organisation
Strategy

Storage needs

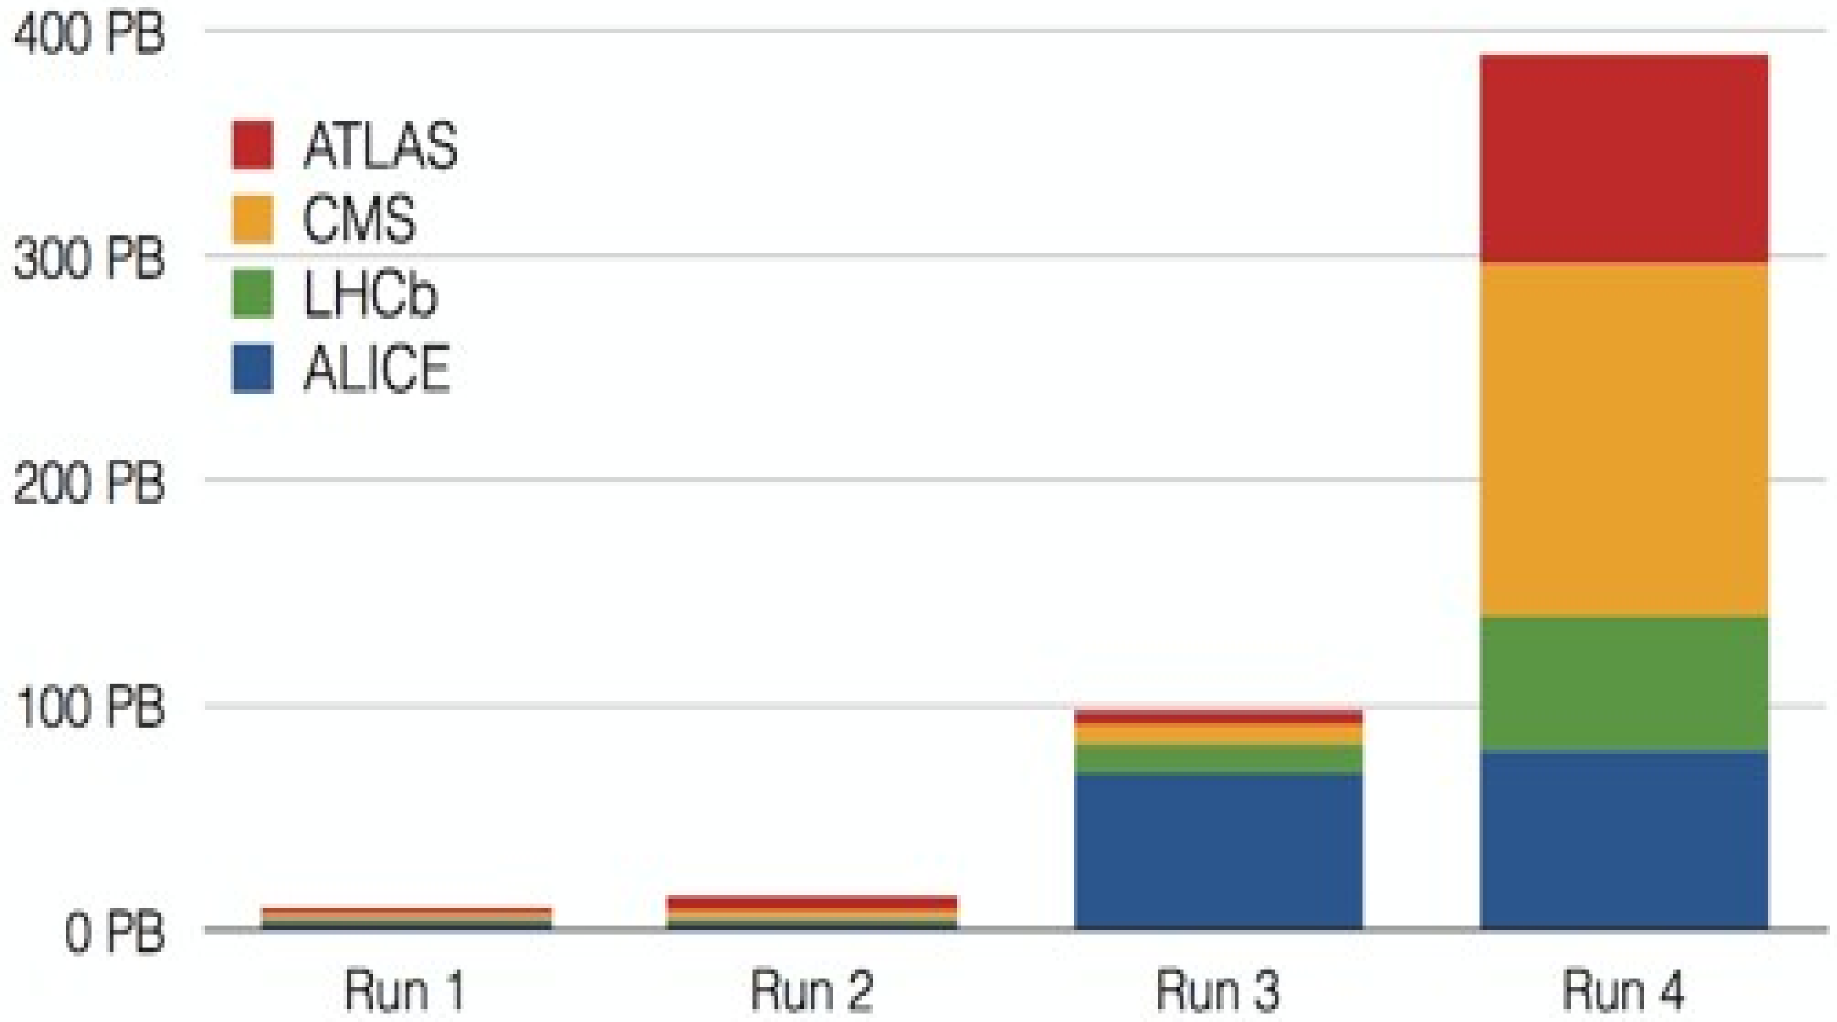

Data per year

	RAW (2 replicas)	Derived	Annual Total	Increase over now
Now	8PB/yr	x8	72PB	x1
HL-LHC do nothing	150PB/yr	x8	1350PB	x18
HL-LHC smart	150PB/yr	x4	750PB	x10

⇨ Fewer replicas
⇨ More network instead

} Up to what point ?

« Network is cheaper than storage »

M. Ernst

Disk cost

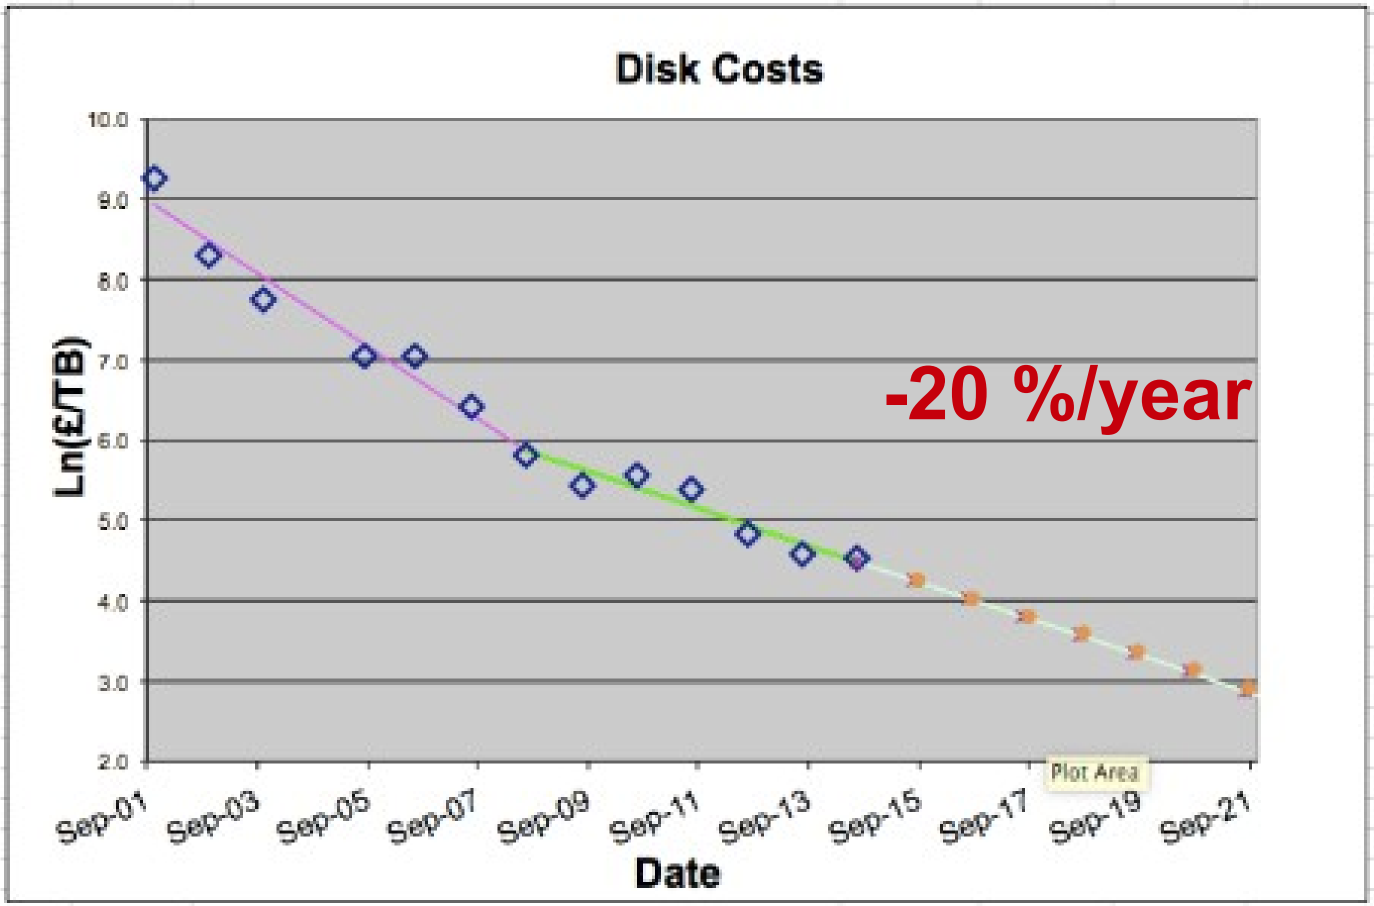

RAL Tier1 Disk Costs with projection

David Britton, Andrew Sansum, GridPP

10 times cheaper in 10 years

... need to be 'smart'

Computing needs

Step	Approx. Fraction Today	HL-LHC do nothing multiplication factor	HL-LHC do nothing CPU increase	HL-LHC smart multiplication factor	HL-LHC smart CPU increase
Generation	0.05	20	1	5	0.25
Simulation	0.45	5	2.25	3	1.35
Digitisation	0.05	20	1	10	0.5
Reco (MC)	0.15	100	15	15	2.25
Reco (Data)	0.1	100	10	25	2.5
Analysis	0.2	10	2	5	1
Total (in units of today's compute)	1		31.25		7.85

Do nothing approach

x 30

Smart approach

x 8

Moore's law

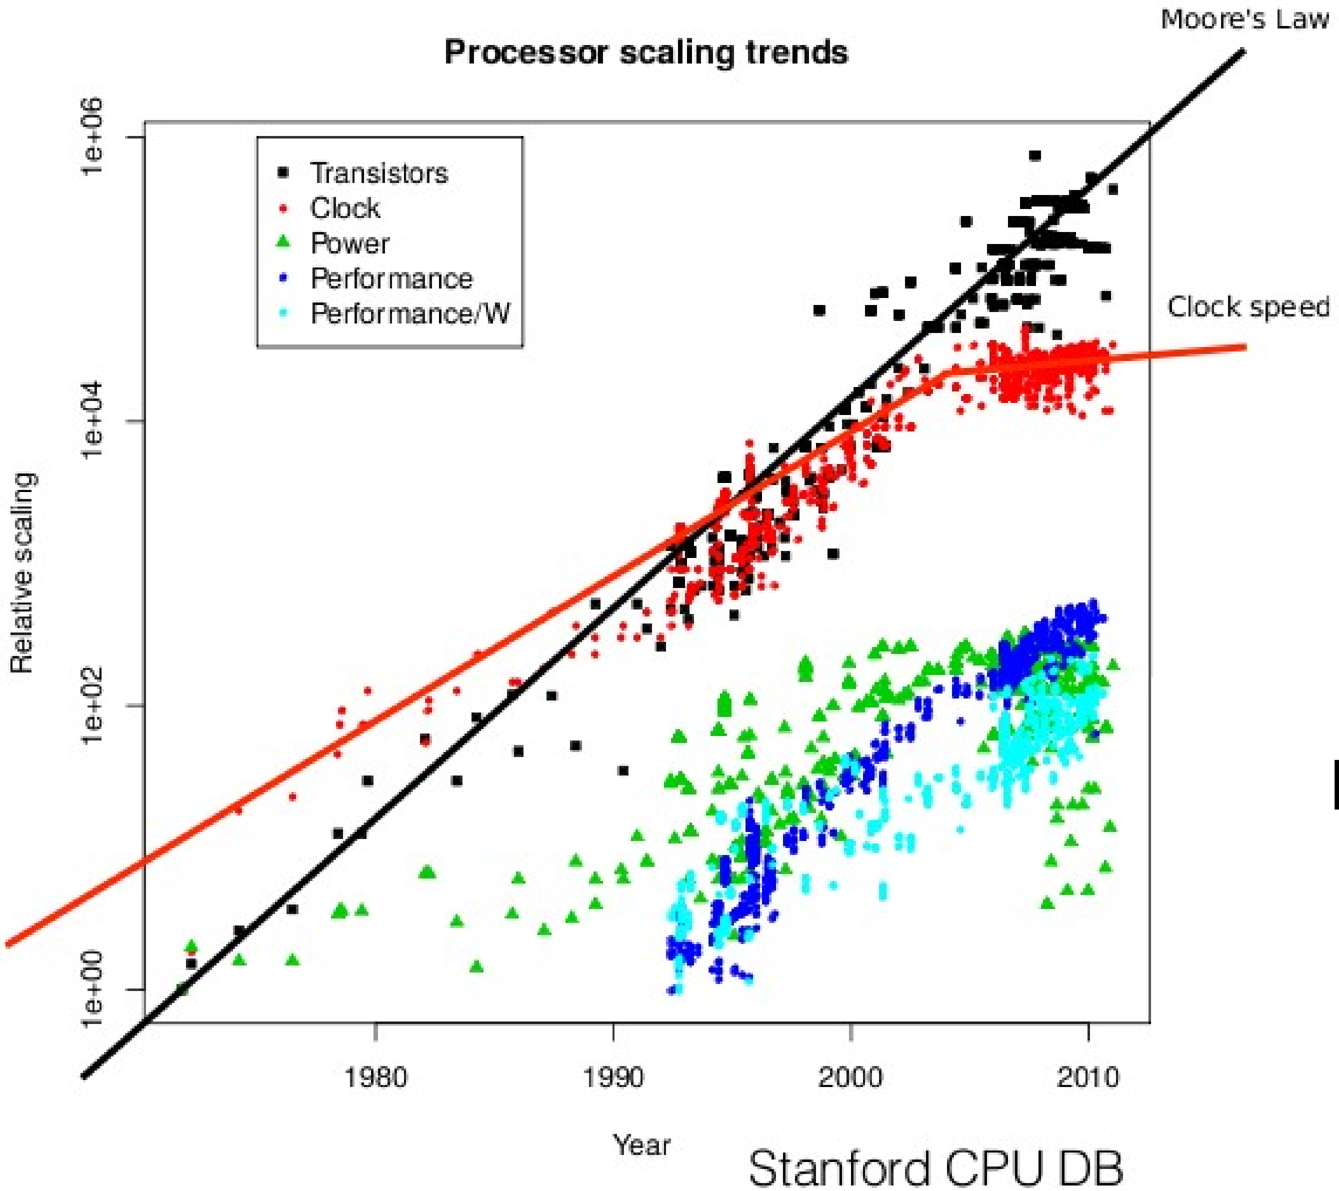

Uncertain!

Frequency ~flat
(won't change)

CPU cost

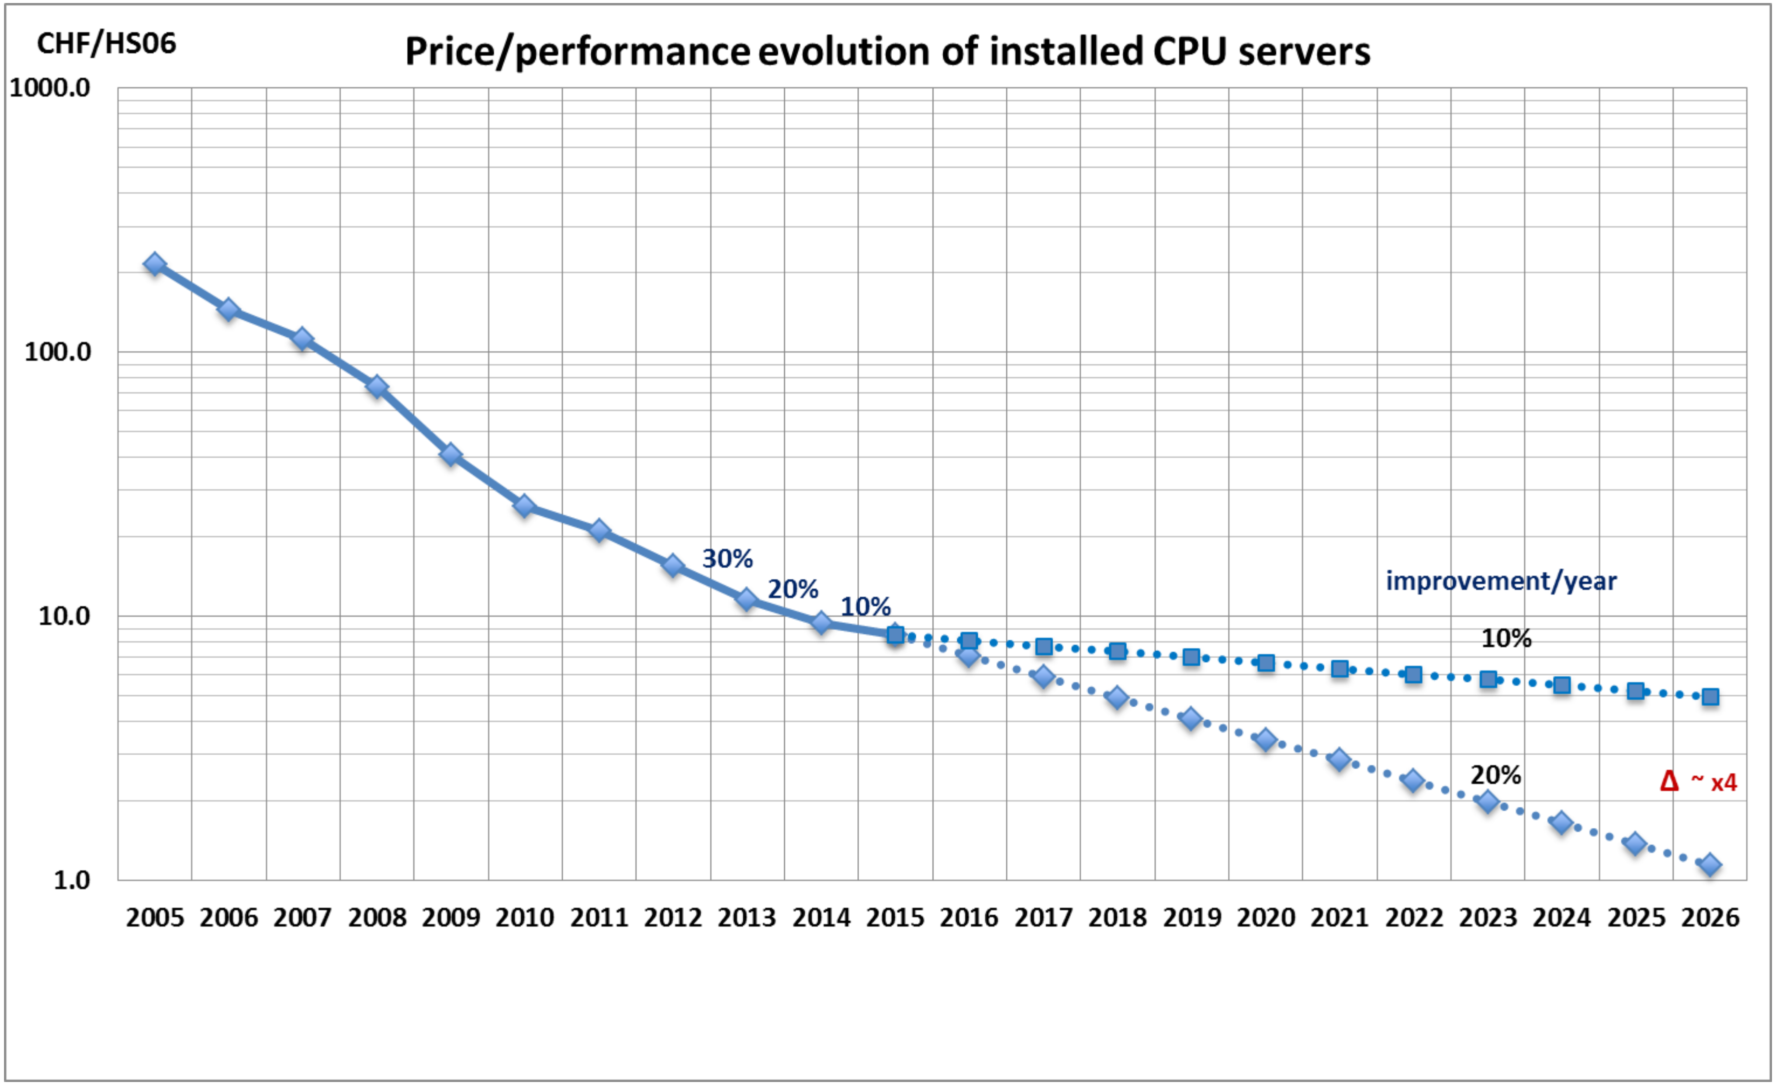

Evolution ??

How to gain CPU ?

- HEP fact
  - CPU functionalities are not used efficiently
- Parallelize
  - Investment
- Fast simulation
  - Quality of physics
- Big penalty « wait for the data »
  - Modern technos not made for that

To which solutions ?

- Design applications based on data structure
- End of Object-Oriented
  - Very convenient for physicists but expensive
- Balance to find
  - Vectorization, memory access, power consumption
  - Maintainability, user-friendliness

Network

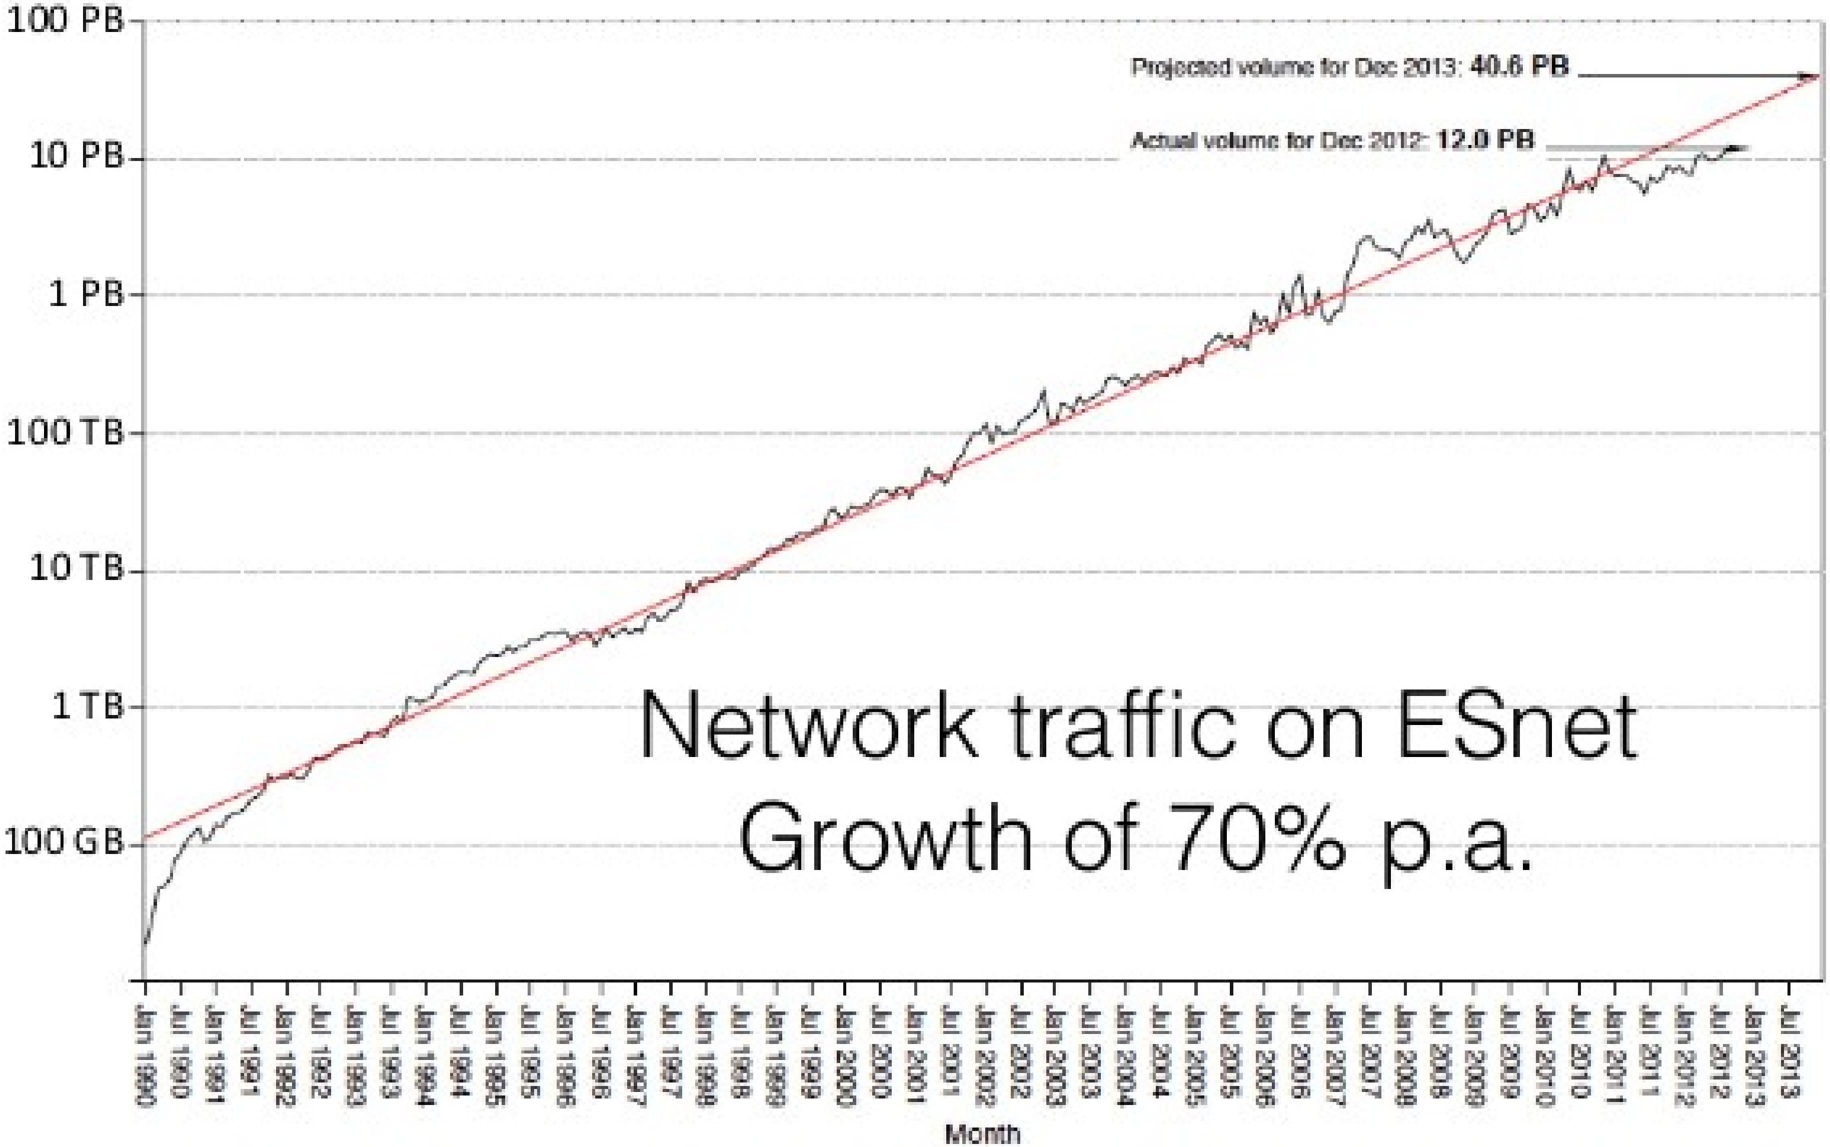

100 PB
10 PB
1 PB
100 TB
10 TB
1 TB
100 GB
Projected volume for Dec 2013: 40.6 PB
Actual volume for Dec 2012: 12.0 PB
Network traffic on ESnet
Growth of 70% p.a.
Month

Performance evolution (flat budget)

Technology	Growth in 10 years	Approche Smart
CPU Servers	x4 - 14	x 8
Disk Capacity	x4 - 10	x 10
Tape Capacity	x10 - 30	x 10
Network Capacity	x30 - 200	cool

Smart approaches should allow us to face'smart' the challenge

Reminder

- HEP not optimal as of today
- But
  - Problem identification and solving
- More and more smart
- Current technologies matures
  - Slower evolution

Diversification ?

- General purpose hardware
  - Little specific problems
- Market law
  - « volume »

Specific hardware
in specific centres
for specific tasks...

Parallelization

- Multithread & multicore
  - ATLAS & CMS & LHCb
- Focus on hot spots
- Simulation : GEANT4

- Gap between needs and post-doc computing skills
- Professional expertise is required

GPU

- Few common and efficient tools
- Parallelize what you need
  - Switch CPU – GPU when needed
- BUT
  - Difficult & costly
- Towards dedicated facilities
  - Cluster-oriented compilation

The memory wall

- Moore ⇨ more & more cores
- Memory per core lowers

- ex. Xeon φ :
  - 60 cores
  - 16 GB memory ⇨ 256 MB/core

⇨ **we must lower the memory footprint**

HPC

- Simulation : 60 %
  - CPU / IO >> 1
- Opportunistic (or not)

- Caveat
  - Experiments must pay the price to make it useful
- HPC
  - Technical & political aspects
  - Who pays ?

Summary

- Still strong scientific motivations for LHC
  - **The** subatomic physics machine
  - Won't have better in the next 20 years
- Computing : « do nothing » approach not viable
  - The landscape is going to change
- Technology evolution
  - Take advantage of it and adapt (ex : HSF)