



Review Calcolo ATLAS Italia

Bologna, 4 settembre 2024

L. Rinaldi & ATLAS-ITALIA computing team

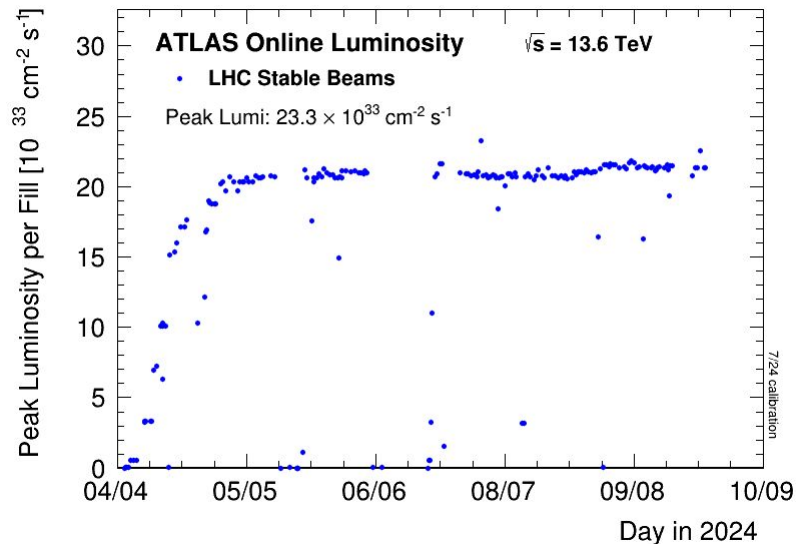

- ▶ Utilizzo risorse calcolo
 - ▶ ATLAS world-wide
 - ▶ Risorse INFN
- ▶ Richieste 2025

Overview ATLAS



Operazioni di presa dati procedono bene

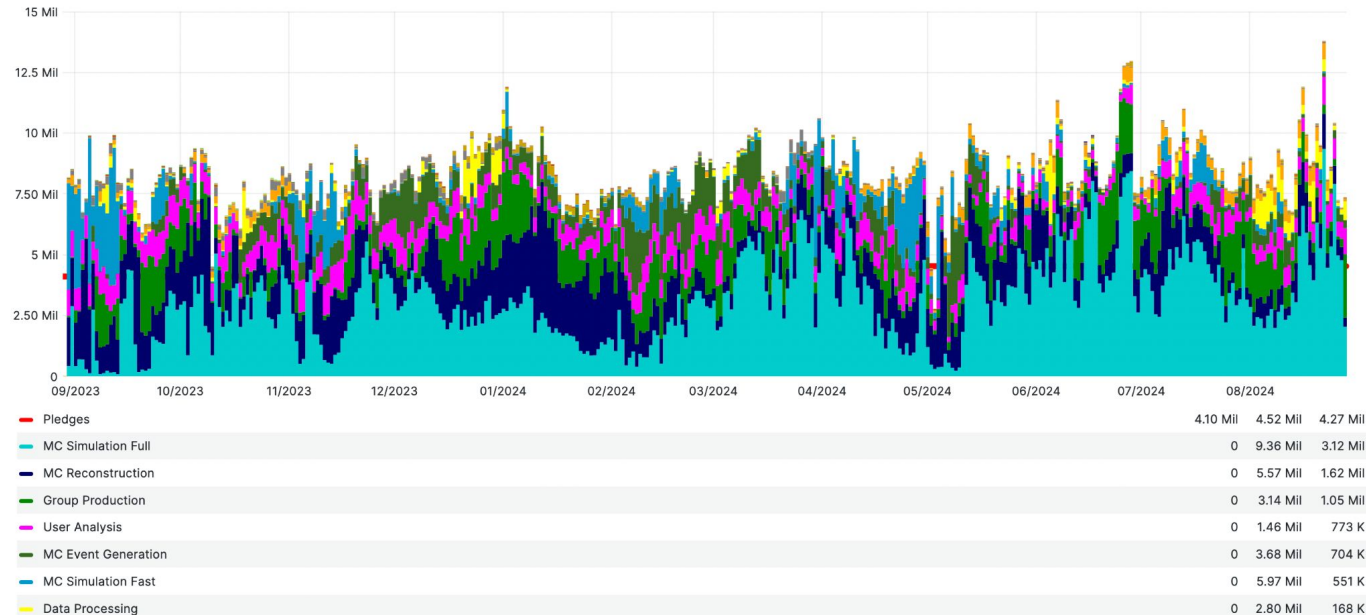
- accumulati 79.2 fb^{-1} (LHC 2024 target 110 fb^{-1})
- Peak Lumi $23.3 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$
- Pileup medio $\mu=64$

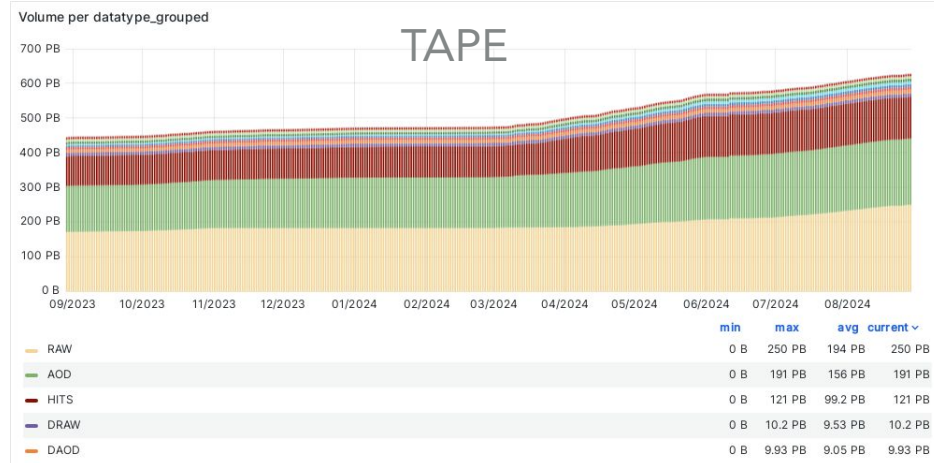
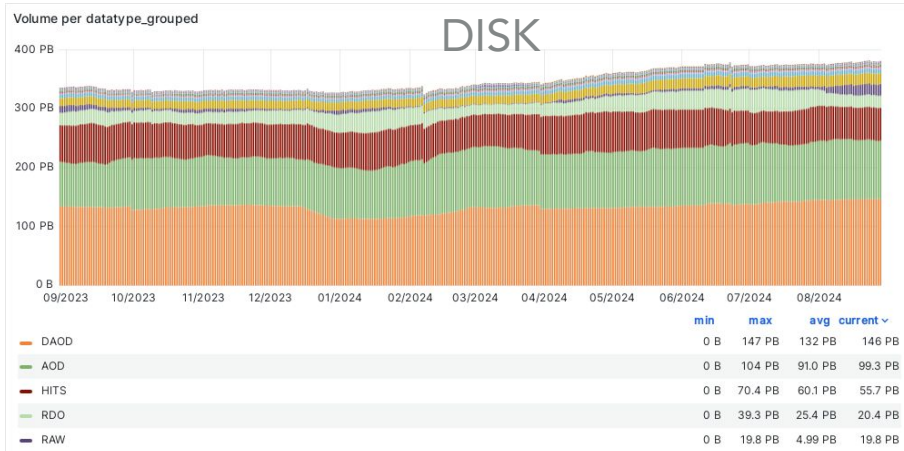


Attività di data processing procedono regolarmente

- dominati da MC-sim, MC-reco, Group-Prod
- in crescita MC FAST simulation

Slots of Running jobs (HS23) ⓘ





Utilizzo dico e tape in costante aumento

- variazione pendenza tape dovuta a presa dati
- + tape carousel model

Overview risorse INFN

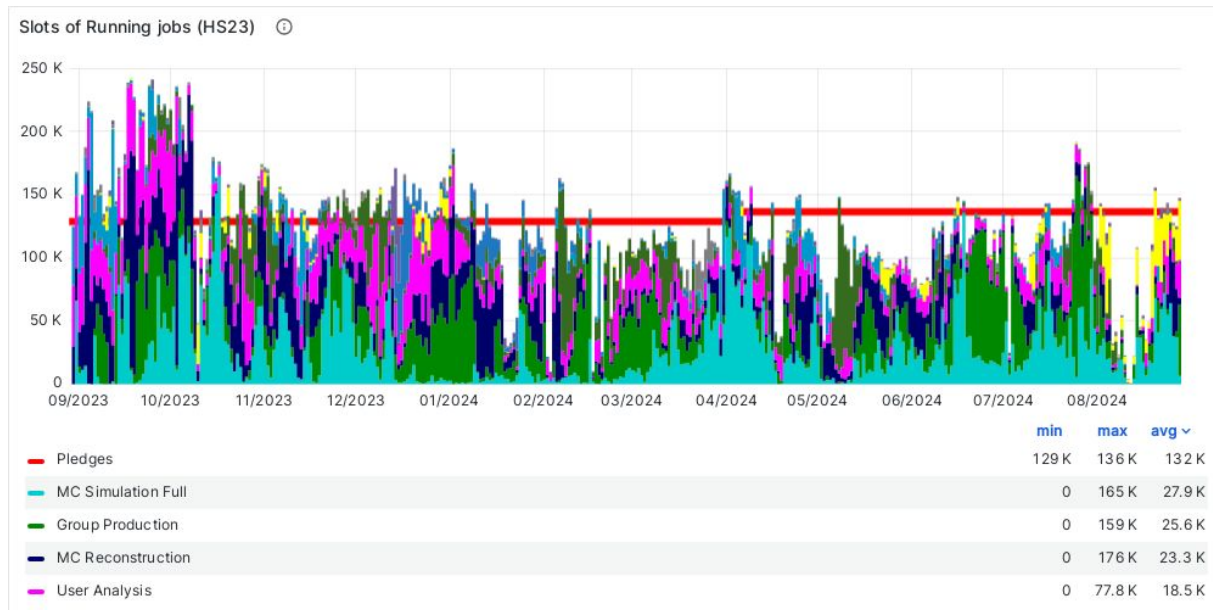
Un po' di sofferenza su farm T1

HS23: media 124 kHS23,
max 242 kHS23

pledge: 136 kHS23

principalmente dovuta a

- generale riduzione della potenza fornita dal T1
- molteplici interventi di sicurezza
- trasloco

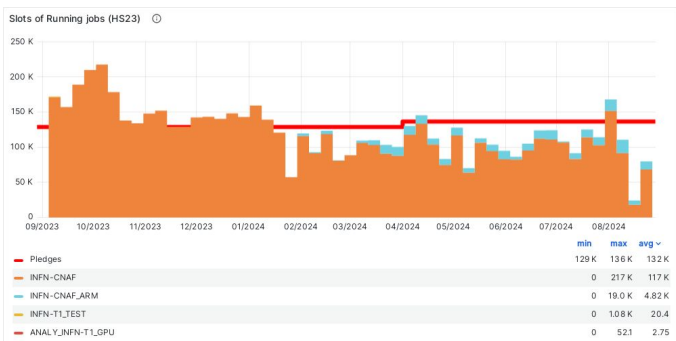


Risorse ARM

4 x Neoverse-N1 + 1 grace
(~17 kHS23, 12% ATLAS-pledge)

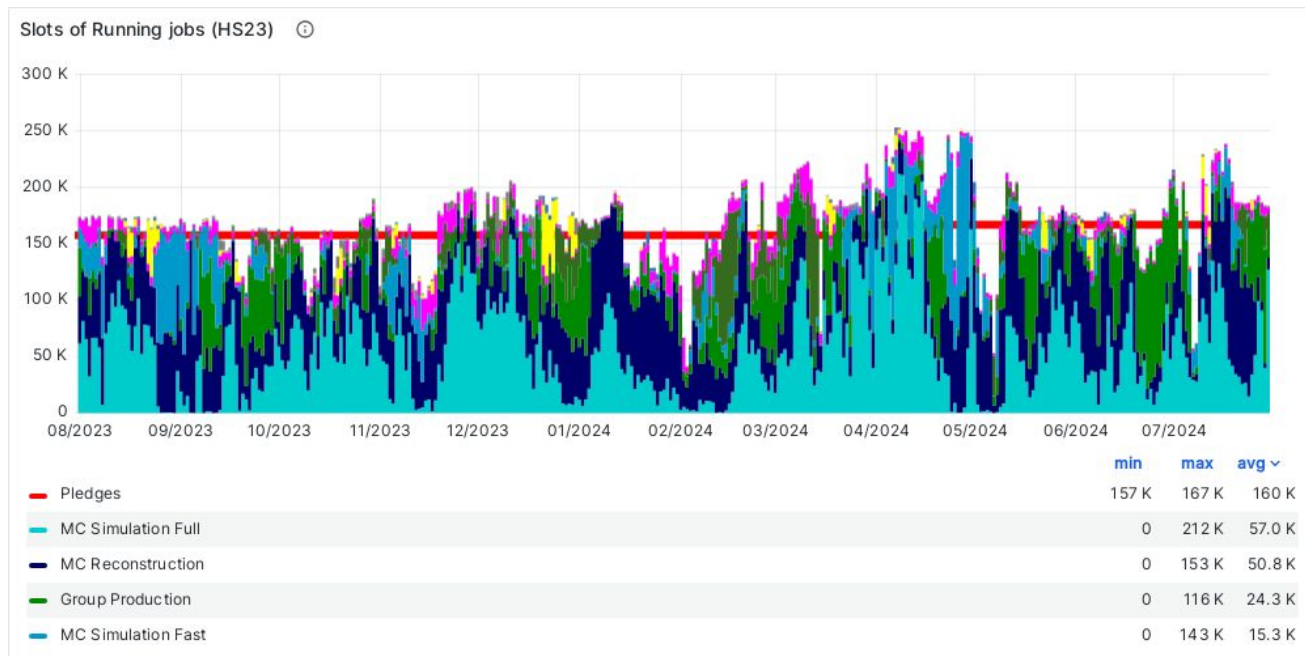
Full validation + utilizzo intenso da parte di ATLAS (MC + group prod)

ARM Flagship (use-case) in CN-HPC Spoke 2



Al momento accounting unico su farm (macchine nello stesso pool)

- da capire come suddividere, anche a valle discussione GDB-giugno (proposta di ARM-pledged)



Complessivamente buona performance dei Tier2

Osservazioni:

ATLAS Dashboard:

value=Ncores*<corepower>

corepower → inserito in CRIC

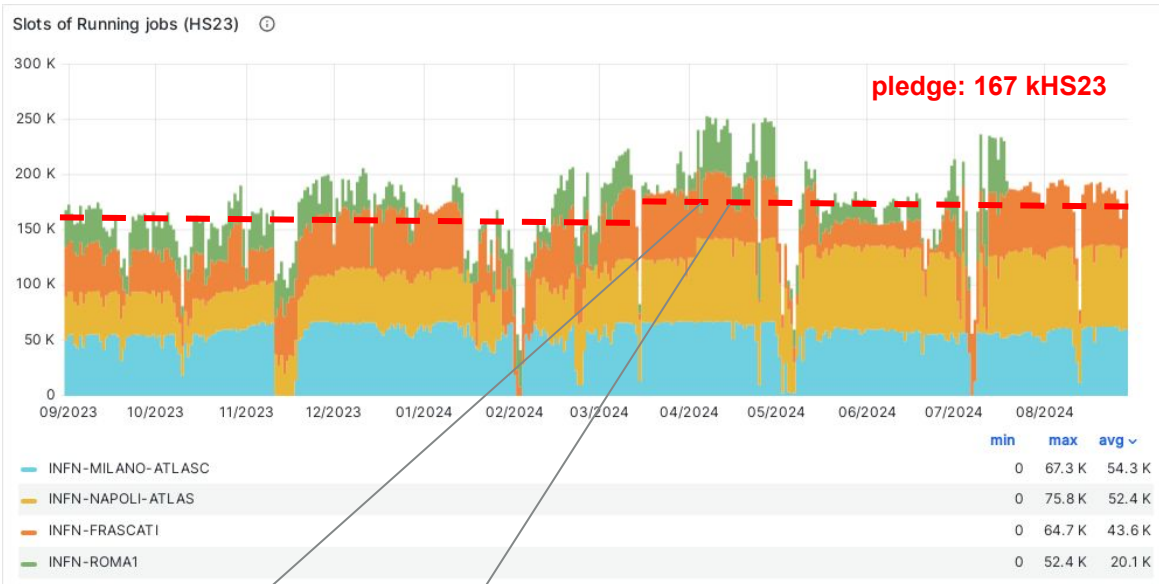
maggio/giugno

- frascati: intervento infrastruttura (lavori sala calcolo)
- roma1: cambiata infrastruttura rete (macchine fuori a blocchi per intervento su switch)

da fine luglio:


- roma1 OFF per power-cut

Tutti i siti riescono a fornire risorse extra-pledge (vecchie macchine...)

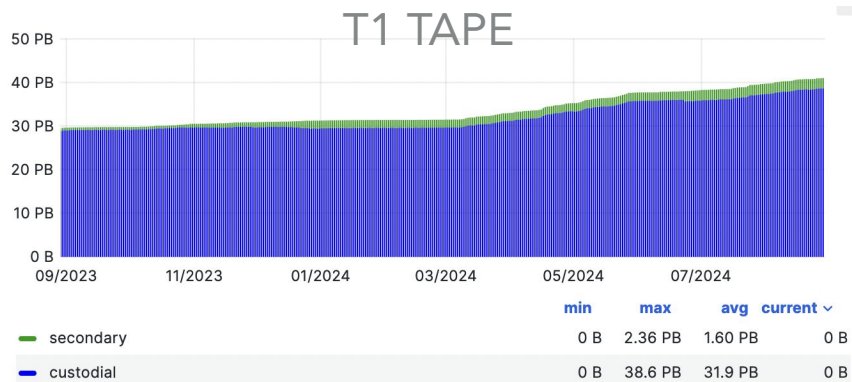
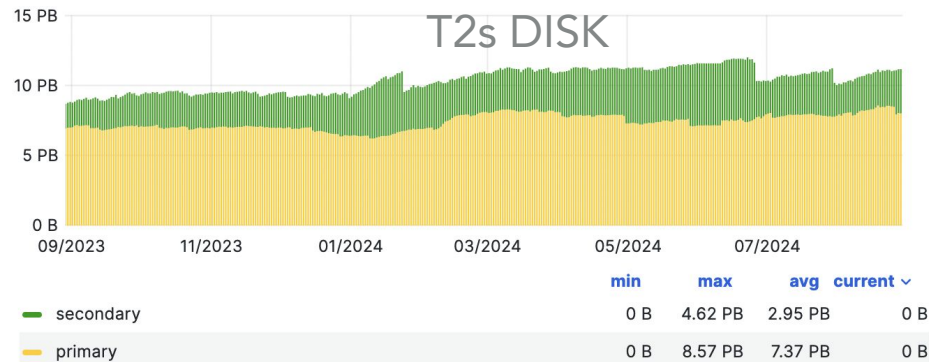



| | |
|--------------------|-----|
| INFN-NAPOLI-ATLAS | 31% |
| INFN-MILANO-ATLASC | 27% |
| INFN-FRASCATI | 24% |
| INFN-ROMA1 | 17% |

share relativo OK in una finestra in cui non ci sono stati problemi



Complessivamente buona performance dei siti italiani



Recente aggiunta di disco al T1
Migrazione a sistema dCache nei
Tier2 (ok Frascati, Napoli, Roma1;
in corso a Milano)

| (dati agosto 2024) | Occupato (PB) | Allocato (PB) | pledge 2024 |
|--------------------|---------------------------------|----------------------------|-------------|
| Frascati | 2.82 | 3.15 | 3.88 |
| Milano | 0.97 (migrazione dCache) | 1.98 | 2.86 |
| Napoli | 5.33 | 5.81 (extra-pledge Ibisco) | 4.81 |
| Roma | 1.87 | 2.18 | 2.62 |
| | | | |
| CNAF-DISK | 13.3 | 13.8(+0.6 buffer) | 14.670 |
| CNAF-TAPE | 41.3 | 40.7 | 40.680 |

Richieste 2025

| ATLAS | | 2023 | | | 2024 | | 2025 | | |
|-------|---------------|------------------|--------------|--------------|------------------|--------------|--------------|---------------------------|------------------|
| | | C-RSG recomm. | Pledged | Used | C-RSG recomm. | Pledged | Request | 2025 req. / 2024 C-RSG | C-RSG recomm. |
| CPU | Tier-0 | 740 | 693 | 675 | 936 | 936 | 1100 | 118% | 1100 |
| | Tier-1 | 1430 | 1462 | 1613 | 1516 | 1514 | 1635 | 108% | 1635 |
| | Tier-2 | 1747 | 1841 | 3457 | 1852 | 2074 | 1998 | 108% | 1998 |
| | HLT | n/a | n/a | 467 | n/a | n/a | n/a | n/a | n/a |
| | Total | 3917 | 3996 | 6212 | 4304 | 4524 | 4733 | 110% | 4733 |
| | <i>Others</i> | | | 1667 | | | | | |
| Disk | Tier-0 | 40.0 | 40.0 | 36.4 | 49.0 | 49.0 | 56.0 | 114% | 56.0 |
| | Tier-1 | 136.0 | 140.9 | 149.0 | 163.0 | 163.1 | 186.0 | 114% | 186.0 |
| | Tier-2 | 168.0 | 160.6 | 155.0 | 200.0 | 194.0 | 227.0 | 114% | 227.0 |
| | Total | 344.0 | 341.5 | 340.4 | 412.0 | 406.1 | 469.0 | 114% | 469.0 |
| Tape | Tier-0 | 174.0 | 174.0 | 127.0 | 207.0 | 207.0 | 258.0 | 125% | 258.0 |
| | Tier-1 | 353.0 | 360.3 | 346.0 | 452.0 | 460.0 | 561.0 | 124% | 561.0 |
| | Total | 527.0 | 534.3 | 473.0 | 659.0 | 667.0 | 819.0 | 124% | 819.0 |

ATLAS C-RSC recommended 2025

| | ATLAS 2025 pledge | Italy share 2025 | ATLASIT 2025 pledge | Richieste 2025 | Richieste 2025 (€) | ATLAS 2024 pledge | Italy share 2024 | ATLASIT 2024 pledge |
|---------------|-------------------|------------------|---------------------|----------------|--------------------|-------------------|------------------|---------------------|
| CPU T1 (HS06) | 1635000 | 0,09 | 147150 | 10710 | 107100 | 1516000 | 0,09 | 136440 |
| Disco T1 (TB) | 186000 | 0,09 | 16740 | 2070 | 207000 | 163000 | 0,09 | 14670 |
| Tape T1 (TB) | 561000 | 0,09 | 50490 | 9810 | 98100 | 452000 | 0,09 | 40680 |
| Totale (k€) | | | | | 412200 | | | |

- ❑ Richieste calcolate rispetto al pledge del T1 (non rispetto alle risorse effettive attualmente disponibili)
- ❑ Mantenuto lo stesso share per ATLAS-IT (9%)
- ❑ Obsolescenza e OverHead **NON** inclusi
- ❑ Fattori di conversione : 100 €/TBn (disco), 10 €/HS23 (cpu), 10 €/TBn (tape)
- ❑ **Allo scorso CdG si è detto che CPU-T1 → Leonardo GP e DiscoT1 → PNRR; resterebbe solo richiesta tape (98100 €)**



ATLAS C-RSG recommended 2025

| | ATLAS 2025 pledge | Italy share 2025 | ATLASIT 2025 pledge | ATLAS 2024 pledge | Italy share 2024 | ATLASIT 2024 pledge |
|---------------|----------------------|---------------------|------------------------|----------------------|---------------------|------------------------|
| CPU T2 (HS06) | 1998000 | 0,09 | 179820 | 1852000 | 0,09 | 166680 |
| Disco T2 (TB) | 227000 | 0,07 | 15890 | 200000 | 0,07 | 14000 |

- ❑ Fattori di conversione : 100 €/TBn (disco), 10 €/HS23 (cpu), 10 €/TBn (tape)
- ❑ Mantenuto lo stesso share ATLAS-IT (9% CPU e 7% disco)

| CPU (HS06) | Frascati | Milano | Napoli | Roma1 | Totale | Pledged |
|----------------------------------------|-----------------|---------------|---------------|--------------|---------------|----------------|
| Totale CPU nei siti per il 2024 (HS06) | 40447 | 41609 | 44546 | 41711 | 168313 | 166680 |
| Obsolescenza 2025 (HS06) | 0 | 0 | 2667 | 0 | 2667 | |
| Delta 2025 (HS06) | 2877 | 2877 | 2877 | 2877 | 11508 | |
| Totale richieste 2025 (HS06) | 2877 | 2877 | 5544 | 2877 | 14175 | |
| Totale CPU nei siti per il 2025 (HS06) | 43324 | 44486 | 47423 | 44588 | 179821 | 179820 |

| | | | | | |
|----------------------|-------|-------|-------|-------|--------|
| Previsione spesa (€) | 28770 | 28770 | 55440 | 28770 | 141750 |
|----------------------|-------|-------|-------|-------|--------|

delta2025=pledge25-totale24

| DISCO (TB) | Frascati | Milano | Napoli | Roma1 | Totale | Pledged |
|----------------------------------------|-----------------|---------------|---------------|--------------|---------------|----------------|
| Totale disco nei siti per il 2024 (TB) | 3877 | 2863 | 4819 | 2619 | 14178 | 14000 |
| Obsolescenza 2025 (TB) | 0 | 0 | 0 | 0 | 0 | |
| Delta 2025 (TB) | 428 | 428 | 428 | 428 | 1712 | |
| Totale richiesta 2025 (TB) | 428 | 428 | 428 | 428 | 1712 | |
| Totale disco nei siti per il 2025 (TB) | 4305 | 3291 | 5247 | 3047 | 15890 | 15890 |

| | | | | | |
|----------------------|-------|-------|-------|-------|--------|
| Previsione spesa (€) | 42800 | 42800 | 42800 | 42800 | 171200 |
|----------------------|-------|-------|-------|-------|--------|

delta2025=pledge25-totale24

| | Frascati | Milano | Napoli | Roma1 | Totale |
|----------------------------------------|-----------------|----------------|-----------------|----------------|-----------------|
| Richiesta CPU (HS23) | 2877 | 2877 | 5544 | 2877 | 14175 |
| Spesa CPU (€) | 28770 | 28770 | 55440 | 28770 | 141750 |
| Richiesta DISCO (TBn) | 428 | 428 | 428 | 428 | 1712 |
| Spesa DISCO (€) | 42800 | 42800 | 42800 | 42800 | 171200 |
| Totale (€) | 71570 | 71570 | 98240 | 71570 | 312950 |
| Overhead server | 5009,9 | 5009,9 | 6876,8 | 5009,9 | 21906,5 |
| Overhead rete | 3866,2 | 3866,2 | 5466,4 | 3866,2 | 17065 |
| Totale (overheads inclusi) (€) | 80446,1 | 80446,1 | 110583,2 | 80446,1 | 351921,5 |

- ❑ Overheads (su CSN1):
 - ❑ quota server : 7% del totale CPU + disco
 - ❑ quota rete : 6% CPU + 5% disco
- ❑ Fattori di conversione : 100 €/TBn (disco), 10 €/HS06 (cpu)

Proiezione obsolescenze CPU (HS23)

| | 2025 | 2026 | 2027 | 2028 |
|---------------|-------------|--------------|--------------|---------------|
| Frascati | 0 | 8440 | 0 | 32000 |
| Milano | | | 14400 | 8400 |
| Napoli | 2667 | | | 72400 |
| Roma | | 14000 | 8400 | 25600 |
| totale | 2667 | 22440 | 22800 | 138400 |

Proiezione obsolescenze disco (TBn)

| | 2025 | 2026 | 2027 | 2028 | 2029+2030 (PNRR) |
|---------------|----------|-------------|----------|-------------|------------------|
| Frascati | 0 | 561 | 0 | 1800 | 2192 |
| Milano | | 850 | | | 2129 |
| Napoli | | 500 | | 3500 | |
| Roma | | | | 873 | 2253 |
| totale | 0 | 1911 | 0 | 6173 | 6574 |

| carica (2024) | nominativo (sez) | kEUR |
|---------------------------------------|------------------|------|
| Resp. Nazionale Calc + Cond DB + VOMS | Rinaldi (BO) | 11,5 |
| computing training | (BO) | 4 |
| Resp Digit | Lari (MI) | 11 |
| SPOT | Marcon (MI) | 11 |

Spese di missione inserite nei preventivi su voce “CALCOLO” esplicitamente




<https://web.infn.it/atlas/organigramma/>

<https://atlaspo.cern.ch/public/ATLASOrganisation/index.html?value=COMPUTING>

Back-up

totale risorse installate nei T2 (con extra-pledge)

| | CPU assegnate ATLAS | extrapledge totale | CPU effettive | corepower cric | max HS atteso | max HS grafana |
|-----------------|------------------------|-----------------------|------------------|-------------------|------------------|-------------------|
| frascati | 40447 | 26269 | 66716 | 11,04 | 736544,64 | 64700 |
| milano | 41609 | 20317 | 61926 | 12,452 | 771102,55 2 | 67300 |
| napoli | 44546 | 30559 | 75105 | 12 | 901260 | 75800 |
| roma1 | 41711 | 15459 | 57170 | 10,81 | 618007,7 | 52400 |
| TOTALE | 168313 | 92604 | 260917 | | | |
| pledge 2024 T2s | 166680 | | | | | |



The Run has been extended by 1 month respect to original plans

- pp@13.5 TeV until Oct 16
- 1 week reference run pp@5.4 TeV
- 2.5 weeks of [Pb-Pb](#)

- LS3 has been extended by 5 months to 3.5 years (closure June 2029)
- However the start of Run-4 is still too tight for ATLAS and CMS
- Proposal under discussion with CERN and FAs: postpone LS3 end to Spring 2030
- Possibly extending Run3 to (part of) 2026
- Decision to be taken by October
- Impact on cost: total impact is zero to first order (CORE does not change, extended operations for Run-3 are compensated by reduced Run-4) anyway there may be a peak in MOF expenditure for 2026 if Run3 is extended.

