

Prescrubbing Meeting Notes

Monday, July 6 2026, noon-3 PM

Attending(Sign yourself in): Shawn McKee, Rob,

Notes by presentation. Please add to the notes. We especially need to track need updates for each presentation in these notes.

WBS 2.3 Overview

Effort should include all FY27, even if it moves to 2.5

Achievement text rewording

Risks need clarification

Granola notes -----

Slides serve as a stage-setter; detailed questions deferred to L3 presentations

Unusual/challenging circumstances called out: significant cost increases assumed to persist

- “Challenging circumstances” was described as “selling it short”
- Slides and subsequent L3 presentations should reflect how this cost reality is being addressed

Org chart updated since last year; all changes marked in red, one change to 2.3.5.3 added same day

- Ivan’s atlas-wide position circled in blue

FTE summary covers FY26 and FY27; numbers pulled from effort spreadsheet and individual presentations

- Discrepancies flagged between slides and effort spreadsheet for 2.3.4 and 2.3.1
- Alexei noted some effort counted comes from people outside 2.3 who were paid for and contributed to 2.3 work; will address in his slides
- Goal: spreadsheet is the truth; slides must follow it

FY27 FTE question: without WBS 2.5, would 2.3 request 27.12 or 29.17 FTEs?

- Answer: 29.17, but acknowledged that a 10% increase is not realistic in the current climate
- One of those FTEs was a new CA position

WBS 2.5 allocation currently zero; all L3s asked to show full FY27 budget in 2.3, with comments on what would move to 2.5

- Ops effort will definitely move; even HI-Ops should stay in 2.3 until formally moved to avoid confusion
- Same approach being taken by WBS 2.2, 2.4, 2.3.5; consistency requested

187 total milestones since FY25; breakdown shown by status including proposed

- Some new 2.3.5 milestones not yet in the table due to missing baseline/end dates

Three new milestones added:

- Facility completes DC27 participation
- Post-DC27 analysis and summary (approx. 6 months after DC27 completes)
- Deploying all pledged resources including end-of-CA purchases (April, medium priority)
 - Priority set to medium because at that point the facility is not below pledge; Paolo pushed back, noting a missed CA purchase would be serious trouble
 - Clarified: milestone is about deployment by that date, not acquisition; six months later still fine

Risk R8-13 (cyber security incident, 2.3 as owner) still not added to risk register despite being proposed last year; Sean still considers it important

New risk suggested by Paulo: robustness and aging of equipment as hardware gets older and creakier

- Sean had been meaning to suggest exactly this risk
- Mitigation options discussed: extending warranties (Southwest Tier 2 already doing this), keeping retired equipment as spare parts, using HPC resources as pledge backfill
- HPC-as-backfill noted as requiring guaranteed use and formal pledge commitment

Accomplishments: US Atlas delivered 22% of all completed jobs, 26% of successful jobs; 25% of analysis jobs; ~25% of storage resources

Challenges reiterated: availability and cost of storage and memory, DC27 prep, EL9 transition, cyber security gaps, personnel difficulties filling positions at competitive salaries

WBS 2.3.2 Tier-2s

Slide 2: FY25->FY26

Slide 5: Add a note about this, including end-of-CA funds

Slide 7: Drop NCE part (offline verify an NCE is not possible)

Paolo: Impact of having LS3 last an additional year?

Personnel essentially unchanged: 9.62 FTEs across the complex, same expected in FY27

- 0.2 FTE of Fin Pang moved to analysis facility effort; replaced by 0.2 of Judith to compensate
- Some people not at 100% because they are in other projects; noted in comments

FY26 funding situation complex:

- Equipment purchase funding level unknown until mid-May; determined to be 1.7M per site
- NSF formally approved; paperwork in place; but NSF has not yet permitted spending
- POs must be encumbered by 2nd November under current rules
- No-cost extension possibility raised by Amir; Fred previously checked with Jono who said no
- Agreed to check again offline; if confirmed no NCE, drop the bullet from slides to avoid raising risks associated with an NCE

Memory costs up 4-6x at peak; average now more or less flat for 16GB DIMMs but was spiking sharply

- Most sites buying large amounts of storage because it ages out faster and is harder to ramp up quickly

Midwest Tier 2 no longer 1.5x larger than other Tier 2s in FY27; targets and pledges now equal for all sites

- Other three Tier 2s see a significant pledge increase as a result
- Midwest Tier 2, already sitting well above 1.5x, now looks “exceptionally rich” by comparison
- Fred flagged this as something that will require careful explanation to Srin

FY27 compute plan: no retirements, no new purchases; run everything until it fails

- End-of-CA hardware included in pledge coverage calculations; needs to be clearly annotated on slides so it is obvious this is not optional

- Verena asked for a summary slide showing what happens to the age profile once end-of-CA funds are folded in, similar to what was shown for UC; to be added

Storage aging: several sites will lose 5+ PB of storage going out of warranty around FY28

- Not visible in FY27 slides; Fred to add the number of storage going out of warranty in FY28
- Future purchases planned on 7-year warranties; if LS3 is extended by one year, this would largely cover the gap
- Most current storage on 5-year warranties; Koshik moving to 6-year; future purchases planned at 7-year

Small task force proposed to think through purchasing strategy given cost environment; considered a wise step

14 of 23 new milestones are M&O activities; effort to better reflect M&O in milestone tracking

- Milestone format: Paulo's requested format used, with one additional line for priority/status level added by Fred; Paulo agreed this was a good idea and should be done across the board

General impression from Paolo: without the cost crisis, 2.3.2 would be performing very well; hardware age profile not worsening significantly; cost plot is the main negative point

WBS 2.3.3 Heterogeneous Integration & Operations

Retire instead of delay milestone 378 on slide 6

Make second milestone on slide 10 "Low" -> "Medium"

Have offline discussion on priority / relevance-to-ATLAS of Globus Compute + Globus + Rucio

Paolo: WBS 2.3.3 never asks for more FTEs. What would you add that would need more effort?

WBS 2.3.4 AFs and Shared Tier-3s

Alexei: We no longer retire things to the AF since costs/availability

Aside: Discussion about getting effort spreadsheet correct and having slides match. Paolo noted that we need to attach milestones to the effort in the DRAFT effort spreadsheet.

Rob: Can adjust slide 3 Hu effort comment to remove Ceph, ServiceX and potentially others?

Slide 3: Paolo: we should call out what we need/want to do, rather than addings lots of “cool” tools/apps we use...too much detail requiring expert understanding.

Slide 4: removed UC Hu line (not a challenge)

Is slide 6 more appropriate for backup?

On slide 8, highlight (via text color) 537/538 which have notes below.

Milestone numbers are messed up in the WBS 2.3 copy spreadsheet. Looks like someone copied and pasted from existing milestones and duplicated existing “Milestone #”s. Duplicated #s are 539, 540, 541, 542

WBS 2.3.5 Continuous Integration & Operations

Slide 2 (Note 2.3.4.5 is NEW?)

Slide 5 (reduce jargon)

Slide 6 (better summary needed reducing detail) Paolo: What is RP-1 ? Rob: The original Kubernetes platform was designed by Lincoln, and now Aidan has pushed this forward as an exemplar that other sites might use. Shawn: Maybe make RP1-core clearer based upon question/response? Paolo: What is Elwood? Ofer: More details later. Paolo: needs tuning for the “management” audience. Rob: I will take a stab at redoing with target audience in mind.

Slide 8: Will need modifications just like 2.3 overview slide: include all effort in 2.3.5.x and then use comments to indicate amount that might move to 2.5 Indicate Shawn’s effort moved FY26 2.3.6: 0.4 to FY27 2.3.6: 0.3

Slide 9: Make sure Ilija’s effort move notes the 0.2 came from 2.4.3 into 2.3.5.2 (reshuffling)

Slide 10: Same mods needed as Slide 8

Slide 11: Do we add some interim effort from Rob and/or Shigeki to bridge from now until the next CA 1.0 FTE hire? There will likely be some period from Feb 1, 2027 until the new hire comes onboard, e.g., if a new hire takes 3 months, the makes available 0.25 FTE. Could we assign that in advance to Rob and/or Shigeki to bridge until then. Maybe Rob at 0.1 FTE?
TBD

Slide 14: Paolo would like to see a closer connection between ATLAS priorities (R2R4 and TDC) effort in 2.3.5.3. Needs (US)ATLAS motivation for why we are doing the listed 2.3.5.3 activities (seem arbitrary)

Ofer will convert slide 21 to the right format.

Slide 22: Lots of jargon; rewrite for a management audience.

Slide 25 Paolo noted WBS 2.3.5.2 and 2.3.5.3 seem to have significant overlap ? Rob: Shannon is a framework to introduce LLM/agents into facility ops in a secure, safe way. 2.3.5.2 is ADC ops-focused Agentic, while 2.3.5.3 develops RP1 as a pattern that can be used broadly at other sites in the future.

---- Granola notes follow:

Four areas: operational support (2.3.5.1), developing services (2.3.5.2), facility R&D (2.3.5.3), cyber security (2.3.5.4, new)

2.3.5.4 implemented without effort allocated in FY26; preparatory work done by Rob and Shigeki

Key FY26 accomplishments:

- Yvon serving as ADC coordinator for the past year; DC27 preparation work ongoing
- Varnish migration for conditions databases completed (described as a big deal)
- ADC analytics platform updated
- Significant agentic AI development by Ilya
- RP1 core deployed: infrastructure-as-code pattern for Kubernetes-based analysis facility
 - Originally a stretched Kubernetes across multiple sites (Lincoln's work); evolved into a deployable pattern for any site
 - Aidan Roseberg (Indiana) made it more accessible and site-agnostic this year
- Open Data Facility (ODF) deployed at UC as proof-of-principle of RP1 core
 - Converged with interest in open data, additional processing capability, and data mirroring picked up in spring
 - Demonstrates that RP1 core pattern can be deployed independently by other sites

Shannon vs. Elwood explained (required ~15 minutes of explanation; Paolo noted this is a problem for scrubbing):

- Shannon: privileged agentic harness for infrastructure management; runs on a dedicated spark with Nemo Claude; controls what AI agents do in a trusted, safe manner at the facility level
 - Motivated by direct experience of Claude breaking things on the AF when given privilege
- Elwood: user-space harness; gives users a trusted set of tools they can run on the AF without facility-level privilege

- Both run on RP1 as the R&D platform; developed and tested there before graduating to production AF and ODF environments
- Security requirements differ between the two; this is the key architectural distinction
- Genesis Phase 1 proposal used to develop the ideas more deeply; very low chance of funding but groundwork laid

2.3.5.2 split into Track A (core services: caching, data services, analytics platform) and Track B (agentic work, candidate for 2.5)

- Track A milestones not yet defined; Rob noted these need to be added

Ilya's effort: moving from 2.4.3 to 2.3.5.2; total goes from 0.8 to 1.0 FTE; 0.6 proposed to move to 2.5

- Net remaining in 2.3.5.2 after move: 0.45 FTE
- Clarified as reshuffling, not new effort; 2.4 slides need to reflect the reduction
- Confusion arose because Ilya's effort was understood to have already moved to 2.3.5.2 last year; spreadsheet shows it was still split 0.8/0.2 between 2.3.5.2 and 2.4.3

Sean's effort: up slightly due to DC27; reclassified from 2.3.2 management (not new total effort)

- Should have been 0.3 in management and 0.2 in 2.3.5.1 in FY26; error not propagated to Chuck's spreadsheet

Ivan's funding: confirmed as US Atlas contributing the full amount; not external Atlas funding
Alec: quarter of ops effort allocated toward AI development work in 2.5; still being discussed with Armin

2.3.5.3 (facility R&D) portfolio described as too jargon-heavy and not tuned for management audience

- Paolo: "run it through an LLM and say, how do I rephrase this for a manager"
- Rob agreed; committed to rewriting slides to be readable for management, removing jargon
- Paolo offered to help; suggested connecting all activities to existing ATLAS R2R milestones and priorities rather than listing tools
- Concern: 2.3.5.3 and 2.3.5.2 appear to overlap in personnel and activities; Shannon/Elwood distinction clarifies the split but is not obvious from slides

Cyber security (2.3.5.4):

- Motivated by: Trusted CI process exposing gaps, August site compromise, and new vulnerabilities introduced by agentic tools
- Work plan developed; liaison model established; OSG security connections made; site security contacts updated; real cyber security exercises begun
- Plan requires 1.0 FTE cyber security lead in next CA; funding source needs verification

- Paolo sees only 0.2 FTE (Rob's) in basis of estimate; proposal may have stated 1.0 FTE; to be resolved Wednesday
- Desired hire profile: someone from the field who can develop tooling, understands the agentic space, and is not a pure compliance officer
 - Rob's slides on the hire profile described as having "nailed it"
- Rob asked whether his effort contribution to 2.3.5.4 ends after scrubbing or continues; answer: fraction of Rob and Shigeki will cover the interim period before the hire arrives
- FY27 1.0 FTE will effectively be split: part to the new hire (once hired), part to Rob and Shigeki in the interim
- Rob to reduce cyber security milestones by factor of ~3; Paolo's guidance: at least 0.2 FTE per milestone; group activities into work-plan-level milestones (~4 total)
- Risk for 2.3.5.4 missing from risk slides: if cyber security lead is not hired, no incident response protocol, no security posture evaluation, sites could be taken offline; should be added or augmented into existing 2.3 risk

Zoom AI Notes from Monday Meeting

Quick recap

This meeting focused on reviewing Work Breakdown Structure (WBS) 2.3, with team members presenting updates on different areas including facility overview, Tier 2 centers, and specific workstreams. Shawn presented the facility overview slides covering accomplishments, FTE summaries, and risk analysis, while highlighting challenging circumstances including significant cost increases for hardware and memory. Frederick presented on Tier 2 centers, noting that Midwest Tier 2 was no longer 1.5 times larger than other Tier 2 centers, which affected pledges and targets for FY27. Rui and Douglas presented on Tier 3.3, discussing workload management, Globus integration, and R&D activities, with discussions about prioritizing IRI/AMSC API integration over Globus Compute. Wei presented on Tier 3.4, covering user monitoring and resource statistics, though some data discrepancies between UC and BNL needed to be addressed. The team discussed hardware cost challenges, potential risks with equipment robustness, and the need to extend warranties as mitigation strategies.

Next steps

Alexei

- Talk with Wei (or others) about the possibility of picking up a fraction of Redwood effort to help with Tania's situation.

Frederick

- Add commentary to the Tier 2 slides to clarify that the end-of-CA funding is not optional and is needed for the hardware purchases.
- Remove the bullet about the possibility of a no-cost extension from the Tier 2 slides, as it may not be possible and could cause confusion.

- Adjust the format of the new milestones in the Tier 2 slides to match the requested format (including priority and status level).

Ofer

- Redo the 2.3.5 slides to make them more understandable for a management audience, reducing jargon and clarifying the main messages.
- Reduce the number of milestones for 2.3.5, consolidating them into fewer, more significant ones.
- Add priorities to the cybersecurity milestones in the spreadsheet.
- Verify and clarify the funding situation for the cybersecurity FTE (whether it is 1.0 or 0.2) before the scrubbing.

Rui

- Retire the milestone related to HPO (no existing use case) instead of marking it as delayed.
- Work with Douglas's to adjust the priority of the milestone related to GPU and IML workflow support to medium (priority 3).
- Adjust the milestone related to Globus Compute API to specify the IRI/AMSC API instead.

Shawn

- Update the FTE summary slide to reflect the correct numbers from the effort spreadsheet, especially for 2.3.1 and 2.3.4.
- Schedule a follow-up meeting for Level 3 (2.3.1 and 2.3.5) on Wednesday at 1 PM.

Wei

- Work on justifying the hardware requests for the analysis facilities by showing resource utilization, age profile, and saturation.
- Consolidate and reduce the number of milestones for 2.3.4, using the standard format and ensuring they are feasible with the available effort.
- Add a milestone for the installation of the end-of-CA hardware for the analysis facilities.

Collaboration

- All Level 3 presenters: Review the shared meeting notes for comments and action items related to their presentations.

Summary

WBS 2.3 Review Meeting

The team reviewed WBS 2.3 overview slides, discussing FTE summaries, accomplishments, and risks. They identified discrepancies between spreadsheet numbers and slide data that need to be addressed. The group discussed new milestones, including participation in DC27,

post-DC27 analysis, and deploying pledged resources. They also talked about risk mitigation strategies, particularly regarding hardware costs and cybersecurity incidents. The conversation ended with plans to move to tier two presentations after addressing some technical issues with sharing access to the presentation materials.

Tier 2 Centers Update Meeting

Frederick presented updates on Tier 2 centers, noting personnel changes including Feng Peng moving from Midwest Tier 2 to the analysis facility and Judith joining Midwest Tier 2 to maintain staffing levels. The team discussed funding challenges, with Frederick explaining that while \$1.7 million in equipment funding was approved for FY26, spending restrictions require purchase orders to be issued by November 2nd. Frederick also highlighted that Midwest Tier 2's relative size to other centers will change in FY27, making targets and pledges equal across all sites.

Hardware Cost and Procurement Challenges

The team discussed hardware cost challenges, particularly rising memory prices that have increased 4-6 times for certain types of memory. They agreed to check with John about the possibility of obtaining a no-cost extension, though Shawn suggested it might be better to drop this option due to associated risks. Frederick presented plans for hardware purchases, noting that most sites are buying significant amounts of storage due to aging hardware, and reported that 2026 milestones were completed well with 14 of 23 new milestones focused on M&O activities. The team also discussed potential implications of a possible one-year extension to long shutdowns, including warranty concerns for storage equipment, and agreed to form a task force to address purchasing challenges.

Area 2.3.3 Progress Review

The team reviewed accomplishments and challenges for area 2.3.3, with Rui presenting updates on automated workload management and FTE allocations. Key challenges discussed included workload efficiency issues and Airdrop failures due to pilot closures, with Douglas noting that ADC tasks might help address the first issue. The team decided to retire a milestone related to Panda inference service due to being proposed too early, while other milestones remain on schedule pending resolution of Globus connection issues to production DCache.

Project Planning and Priorities Discussion

The team discussed plans and priorities for their project, focusing on finishing the NERSC allocation and improving efficiency. They debated the priority level of integrating the IRI/AMSC API instead of Globus Compute, ultimately agreeing to keep it as medium priority. Alexei raised concerns about resource allocation and alignment with Atlas needs, while Ivan confirmed RUSHO's interest in the Globus plugin implementation. The team reviewed milestones based on their plans, with five new milestones created, and discussed the potential for scaling up if additional resources became available.

Resource Utilization Data Discrepancies

Wei presented monitoring data from UC facilities showing user counts and resource utilization, but identified discrepancies in the BNL database numbers that need debugging. Paolo emphasized the importance of demonstrating resource constraints at Slack and Brookhaven facilities, suggesting Wei should prepare utilization data to justify continued operations despite potential user declines. Shawn highlighted discrepancies between the slides and effort

spreadsheet numbers, stressing the need to align all documentation before finalizing reports, as the spreadsheet serves as the basis for funding allocation.

Resource Constraints and Support Challenges

The team discussed resource constraints and support challenges across three facilities, with Rob explaining that limited staff and infrastructure make it difficult to implement cross-facility initiatives like unified login systems. Paolo suggested focusing on describing intended outcomes rather than listing specific tools, and the group agreed to rebalance the presentation by removing some items and adding more context about tool purposes. The discussion highlighted ongoing GPU demand issues, particularly for LLM access, with Paolo noting that providing sufficient GPUs would require significant additional resources beyond current capabilities.

Budget and AI Access Planning

The team discussed budget allocations, with Rob explaining that the 71K funds came from the Midwest Tier 2 budget and AF budget. Wei presented ongoing activities and milestone reviews, noting that two technical milestones (537 and 538) needed to be retired due to personnel changes and lack of user interest, though he plans to propose similar initiatives in FY27 using different technology approaches. The discussion focused heavily on AI/LLM access needs, with Wei proposing solutions for graduate students and exploring options between commercial services, NRP's open models, and potential use of MSAC HPC resources, though no final decisions were made on the specific approach.

Hardware Replacement and Milestone Planning

Wei presented hardware replacement numbers for Slack, but Paolo requested more detailed justification showing current hardware saturation and age profiles before approval. The team discussed 27 planned technical milestones for 2027, with Wei proposing to shrink the remote access scope to focus on local group disk space allocation and demonstrating quick data movement capabilities. Paolo questioned whether these activities should be classified as operational or R&D milestones, noting a contradiction with previous statements about limited effort allocation.

Milestone Feasibility and Formatting Issues

The team discussed concerns about milestone feasibility and formatting issues in their presentation slides. Paolo expressed worry that the current milestone structure, including dates as far as September 2027, appears unfeasible and lacks clear ownership. The group identified duplicate milestone numbers in the spreadsheet that need to be resolved, and Verena requested the addition of missing items like UCAF hardware installation and end-of-CA information. Paolo emphasized that the current slides are not ready for the upcoming scrubbing meeting, warning that the reactions could be negative, and requested significant improvements to the presentation and effort spreadsheets within the next week.

Task 231 and Facility Updates

The team discussed rescheduling work on tasks 235 and 231, deciding to focus on task 231 on Wednesday at 1 PM instead of completing both tasks today. They agreed to replace the regular facility coordination meeting with the task 231 discussion. Ofer presented an overview of activities showing four areas including operational support, developing services, facility R&D, and Level 4 work, highlighting accomplishments such as Yvonne's role as ADC coordinator, database migration to Varnish, and Agentic AI development. The team clarified details about

RP1 (Research Platform 1) and its deployment as an open data facility, with Rob explaining it was designed to be used by multiple ATLAS sites rather than being specific to their analysis facility.

Presentation and Resource Planning Updates

The team discussed revisions needed for presentation slides to make them more understandable for management, with Rob agreeing to redo the slides to remove jargon. They reviewed and adjusted resource allocation tables for different work areas, including moving some effort from FY26 to FY27 and clarifying Ilya's effort distribution between different project components. The discussion concluded with debate about interim staffing needs for the facility R&D role, with agreement that the planned 1.0 FTE would need to be split between Rob and any future hire during the recruitment period.

Cybersecurity Hiring and Budget Planning

The team discussed cybersecurity hiring plans, with confusion arising about whether 0.2 or a full FTE was budgeted for the position. Rob explained the distinction between Shannon (a privileged system for infrastructure management) and Elwood (a user-space version for physicist analysis), both part of a broader Agentic AI development framework. The group reviewed ongoing activities and challenges, including operations during the long shutdown period and preparation for DC27 in February-March. A discrepancy was identified between the proposed and actual budget for cybersecurity staffing, which needs to be clarified with the NSF.

Cybersecurity Milestones and Funding

The team discussed cybersecurity milestones and funding requirements, with Paolo advising that milestones should not be added to the spreadsheet until priorities are clarified and funding sources are confirmed. Rob agreed to reduce the number of milestones by a factor of three, and the team noted that presentation time allocations for upcoming scrubbing sessions may need adjustment to accommodate the current slide volumes. The discussion concluded with plans for additional review sessions on Friday between presenters to refine their messages and address outstanding questions about facility distribution and R&D priorities.

US ATLAS Pre-Scrubbing Meeting — WBS 2.3 Facilities and Distributed Computing

Meeting Metadata

- **Title:** US ATLAS Pre-Scrubbing (WBS 2.3 Facilities and Distributed Computing)
- **Date:** Monday, July 6, 2026
- **Recorded participant (Granola account):** Robert Gardner
- **Presenters/speakers identified in this session:**

- Shawn McKee (University of Michigan) — WBS 2.3 Subsystem Manager; chaired the meeting and presented the 2.3 Overview
- Alexei Klimentov (BNL) — WBS 2.3 co-PI, WBS 2.3.1 Tier-1 lead
- Fred Luehring (Indiana University) — WBS 2.3.2 Tier-2 co-lead, presented 2.3.2 (Rafael Coelho Lopes de Sa, the other 2.3.2 co-lead, was unable to join — traveling in India)
- Rui Wang (Argonne National Laboratory) — WBS 2.3.3 lead, presented 2.3.3
- Doug Benjamin (BNL) — co-presenter for WBS 2.3.3
- Wei Yang (SLAC) — WBS 2.3.4 lead, presented 2.3.4 (Analysis Facilities: BNL AF, SLAC AF, UC AF)
- Ofer Rind (BNL) — WBS 2.3.5 (Continuous Integration and Operations) lead, presented 2.3.5; co-authored with Ivan Glushkov (BNL)
- Robert Gardner (UChicago) — UC AF PI and WBS 2.3.5.3/2.3.5.4 contributor; an active participant throughout, and from the 2.3.4 presentation onward his turns are explicitly labeled in the source transcript
- Ivan Glushkov — ATLAS ADC Coordinator (BNL); commented on Globus/Harvester integration during 2.3.3 and co-authored the 2.3.5 deck
- "Paolo" and "Verena" — asked most of the review questions throughout, with Paolo consistently the one urging presenters to move through slides rather than read them verbatim; full names/affiliations were not stated in the recording
- **Agenda covered in this transcript:** Administrative opening; WBS 2.3 Overview (Shawn McKee); WBS 2.3.2 Tier-2 Centers (Fred Luehring); WBS 2.3.3 Heterogeneous Integration & Operations (Rui Wang / Doug Benjamin); WBS 2.3.4 Analysis Facilities (Wei Yang); WBS 2.3.5 Continuous Integration and Operations (Ofer Rind). The session ran out of time for WBS 2.3.1 (Tier-1).
- **Agenda note:** Alexei Klimentov had originally planned to leave early, so Tier-1 (2.3.1) was moved to the end of the agenda; he ended up staying through both 2.3.4 and 2.3.5. With no time left for 2.3.1, the group moved it to Wednesday, July 8, 2026 at 1:00 pm (replacing that day's regular facility coordination meeting), with Alexei presenting. A separate working session between Alexei Klimentov and Rob Gardner to rework the 2.3.4/2.3.5 slides was tentatively proposed for Friday, July 10, 2026, ahead of the July 13–17 scrubbing.

Reference Links

- 2026 WBS 2.3 Scrubbing Google Drive folder (all meeting materials):
<https://drive.google.com/drive/folders/1zIYxY4WDNtbMQCI9uCgNaZwuCTOLMnjf>
- WBS 2.3 2026 Overview Scrubbing slides (Shawn McKee):
https://docs.google.com/presentation/d/1VPkjinSuchkYPVw1IToc78DN8L8AEbr-EJRU_DvttmaY/edit
- WBS 2.3.2 2026 Scrubbing slides (Fred Luehring / Rafael Coelho Lopes de Sa):
<https://docs.google.com/presentation/d/1EC9GNy8I93zoiYI23uBIgaZzDB3WRHSJoIXjZOdiL9g/edit>

- WBS 2.3.3 2026 Scrubbing slides (Rui Wang):
<https://docs.google.com/presentation/d/1HFXKglkqra1v3CVOslevYg1BAEU3DqqBZi7hb pXo1OA/edit>
- WBS 2.3.4 2026 Scrubbing slides (Wei Yang):
https://docs.google.com/presentation/d/1Ko_0NSrXbbkHK6TidVEQ8vN13h6vXIHD2-AX EhyLtf0/edit
- WBS 2.3.5 2026 Scrubbing slides (Ofer Rind / Ivan Glushkov):
<https://docs.google.com/presentation/d/1KOuQoCvxS4cAqB7fzY94Sa9VNisPtbPymF3L yHHa1mE/edit>
- WBS 2.3 milestone tracking spreadsheet (referenced in meeting as "the copy of the milestones for WBS 2.3"):
https://docs.google.com/spreadsheets/d/1FkVDqLh_5PaHQDP-bfefBJ-PlolfD7LLw3sbP _vTgB0/edit
- Scrubbing Indico page (for the July 13–17, 2026 in-person scrubbing) and the shared meeting-notes document were both referenced verbally in the recording; their URLs were not read aloud and are not captured in this transcript. Both should be in the Drive folder linked above (Indico for today's call also carried these links, per Shawn's opening remarks).

Summary

This was a prescrubbing walkthrough ahead of the July 13–17, 2026 WBS 2.3 scrubbing review, run as a slide-by-slide read of the Overview, Tier-2 (2.3.2), Heterogeneous Integration & Operations (2.3.3), Analysis Facilities (2.3.4), and Continuous Integration and Operations (2.3.5) decks with the review committee (Paolo and Verena asking most questions and Paolo consistently pushing pace, Alexei Klimentov weighing in throughout on effort accounting, HPC/Globus strategy, and AF/facility strategy). The session ran out of time for WBS 2.3.1 (Tier-1), which was moved to Wednesday, July 8. The meeting closed with a blunt overall assessment from Alexei that the 2.3.4 and 2.3.5 material is not yet ready for the real scrubbing.

2.3 Overview (Shawn McKee). Shawn framed the year around one dominant theme: an unprecedented hardware cost environment (a reviewer pushed back that "challenging circumstances" undersold this) compounding a \$1M FY26 budget cut at BNL and continuing difficulty finding and retaining skilled staff at competitive salaries. The FTE summary generated the most discussion: WBS 2.3 totals 27.12 FTE in FY27 versus 26.97 in FY26, but that FY27 number already assumes roughly 2.05 FTE moving to the new WBS 2.5 (AI/ML) area; without that move, FY27 would be 29.17 FTE. Alexei flagged that some 2.3.1 and 2.3.4 effort numbers on the slide don't match the source effort spreadsheet because they include people outside 2.3 whose effort is paid from 2.3 sub-WBS budgets — he'll clarify this in his own slides. On FY27 budget mechanics, the guidance to all Level 2 areas (WBS 2.5 currently carries a zero allocation) is to submit a full FY27 budget under 2.3, flag in comments what should eventually move to 2.5, and wait for a coordinated, explicit move rather than pre-splitting the budget now. On delivery, US ATLAS accounted for roughly 22–26% of global Tier-1/Tier-2 completed and successful jobs and about 25% of storage and analysis-job share. 187 milestones were tracked

across 2.3.x since FY25 (some new 2.3.5 milestones are still pending baseline dates). On risk, Shawn will formally add R8-13 (a facility cybersecurity incident, with 2.3 as risk owner) to the register — proposed last year but never entered — and, prompted by Paolo, will add a new risk around aging/robustness of facility hardware given the extended-use mitigation strategy the facility is now relying on. Three new FY27 milestones were added: completing DC27 participation, completing post-DC27 analysis/summary work (~6 months after DC27), and deploying all pledged resources including end-of-CA purchases (medium priority, since the facility is not below pledge in the interim).

2.3.2 Tier-2 Centers (Fred Luehring). Total Tier-2 effort holds steady at 9.62 FTE across AGLT2, MWT2, NET2, and SWT2 for both FY26 and FY27 (Fengping Hu's effort shifted from MWT2 to the Analysis Facility, offset by bringing on 0.2 FTE of Judith Stephen). Hardware-wise, all four sites are meeting or exceeding FY26 compute and storage targets, and the fraction of very old hardware (>8 years compute, >6 years storage) is small and shrinking, aided by heavy investment in storage refresh. FY27 introduces a structural change: Tier-2 pledges and targets become equal across sites, ending MWT2's historical 1.5x-larger allocation — this significantly increases pledges at AGLT2, NET2, and SWT2, while making MWT2 look comparatively over-resourced, a framing Fred wants to get ahead of before it's raised externally. FY26 was complicated by two forces: the \$1.7M-per-site equipment funding level wasn't confirmed until mid-May, and NSF has approved spending flexibility but not yet released funds, with a hard requirement to encumber (issue POs) by November 2, 2026 — against a hardware market where memory/SSD prices have risen roughly 4x and HDD prices roughly 2x. A no-cost-extension option was raised, but Fred has already checked with Jono and been told no; the group agreed to recheck once more and drop that bullet if it's confirmed unavailable, partly because pursuing it may carry its own risk. Milestones: 2026 milestones mostly closed, with only three minor delays; 14 of 23 new milestones are M&O-related, and Fred added an explicit priority/status field to the backup slides at Paolo's request (a change other L3 areas should probably standardize going forward). Looking to FY27, the first part of the year will be dominated by procuring and installing roughly \$2M in new equipment per site, timed to also support DC27 mini-challenge preparation. A possible one-year extension of Long Shutdown 3 was discussed as a mixed bag for Tier-2s — likely net positive (less pressure to ramp up quickly) but it would push more of the newly purchased storage past warranty; the group floated pre-paying for extended (7-year) warranties on the current storage buy as a hedge, and agreed a small task force may be worth forming to think through the purchasing/warranty tradeoffs.

2.3.3 Heterogeneous Integration & Operations (Rui Wang / Doug Benjamin). FY26 accomplishments centered on integrating Globus Compute with the NERSC Superfacility API and Harvester (with HEP-CCE) and demonstrating Rucio+Globus data-transfer capability at NERSC. Combined 2.3.3/2.4.2 effort is 0.7 FTE in both FY26 and FY27, proposed to move to WBS 2.5. The most consequential open issue is that Globus has not yet been reconnected to BNL's production dCache, so HPC sites currently move data by having pilots copy directly to/from BNL dCache rather than through the officially recommended Globus path — workable for now (NERSC hasn't objected) but not sustainable into FY27 without the reconnection. Alexei Klimentov pressed on this and separately raised a broader concern: with only 0.7 FTE, the team can't chase every HPC integration path (Globus Compute, Harvester, the evolving IRI API)

without ATLAS ADC/data-handling experts confirming the direction has real value to the collaboration; Rui noted the data-handling group itself had asked for a status talk at the last collaboration meeting and there were no strong objections raised there. Ivan, speaking for ADC/PanDA, confirmed the team would welcome a Globus plugin for Harvester developed by Rui. Alexei said he remains unconvinced and asked to take the Globus-vs-IRI-API strategy discussion offline. Milestone housekeeping: several HPC allocation milestones tied to the retired Frontera system were retired rather than marked "delayed," and the group agreed retiring (with the option to reopen later) is the right status when there's genuinely no current use case, rather than "delayed," which implies a missed commitment. On effort, when Paolo asked directly why the team never asks for more FTEs, Doug Benjamin's answer was that scope should be set at the WBS 2.5 level given funding constraints, and that HEP-CCE collaboration lets the team leverage outside effort rather than needing to grow headcount — though he acknowledged there are things (e.g., a broader Globus overlay-cluster approach, which Paolo noted CMS already uses) the team isn't pursuing simply because of the 0.7 FTE ceiling, and any such expansion would need ADC/Zach's explicit buy-in.

2.3.4 Analysis Facilities (Wei Yang). This section ran noticeably rougher than the first three, and several of the problems raised need Rob's own follow-up before the real scrubbing, not just Wei's. Accomplishments opened with usage stats across the three AFs (SSH/Jupyter/batch users and batch jobs/hours), but Wei flagged the BNL numbers as unreliable — they look "completely off" compared to the other two AFs, and a re-pull the day before the meeting gave different figures than the June 11 snapshot on the slide; he's waiting on Qiulan Huang to debug the BNL AF monitoring database. Paolo pushed hard on the narrative around declining SLAC user counts, essentially previewing the question a real reviewer would ask ("why keep operating SLAC procurement if usage is falling?") and asked Wei and BNL to come prepared with utilization plots showing that their smaller, resource-constrained footprints are actually saturated, framed positively rather than as a bare "we're constrained" complaint. The FTE summary surfaced two concrete data problems: (1) the slide's FY26→FY27 total (1.05 → 1.25) doesn't reconcile — Chicago's Hu goes up 0.20 and BNL's Pelosi is newly added at 0.10, which by Alexei's on-the-spot math should push the FY27 total to 1.35, not 1.25; and (2) the effort-spreadsheet-vs-slide mismatch flagged in the Overview recurred here: Nakata appears on the FY26 slide but isn't in the master effort spreadsheet for FY26, and the person Nakata replaced still shows under their old name in Chuck's spreadsheet. Alexei separately noted BNL's new 0.1 FTE for Pelosi was never actually discussed with BNL. Paolo also questioned why the Chicago FTE portfolio (Kubernetes, notebooks, Dask, HTCondor, Ceph, BinderHub-Fed, IaaS, REANA, AI-ops, ServiceX) looks so much larger and more scattered than the other two AFs' lines, and — echoing a critique he'd already made on the 2.3.5 slides — argued the slide should explain *why* each effort exists, not just list tool names. Rob agreed and noted some of that portfolio (the BinderHub/federated-K8s/notebooks/Dask/HTCondor line, and ServiceX, which is gated by and already counted under 2.3.5.2) could reasonably come out of the 2.3.4 table, since Chicago's AF has always leaned on Midwest Tier 2 infrastructure support beyond its own 0.3 FTE base. On hardware, Paolo declined to accept a flat replacement-cost ask and told Wei/BNL to first show current hardware is saturated, then show its age profile, then justify a multi-year buy plan — homework that still needs doing before the FY27 ask (\$331,451

for UC AF alone, per the slide, plus per-site requests at BNL and SLAC) is presented for real. The milestone review surfaced its own housekeeping problems: milestones 539–542 appear to be duplicates of pre-existing entries that Wei can't trace the origin of; several 2.3.4 milestones use "ongoing" or "quarterly" in the date fields, which Shawn McKee flagged as inconsistent with how recurring milestones are tracked elsewhere (Tier 2's fixed-period, reissue-next-period pattern); and the deck mixes milestone-table formats, which Verena asked to standardize. The biggest substantive dispute was Local Group Disk: Wei's slide, built on a WBS 5 user survey, treats LGD as underutilized and proposes reorganizing/shrinking it; Rob, drawing on roughly 20 RAC meetings over more than two years, said he has never once heard a user complain about LGD shortage and doesn't know where the milestone's premise comes from — this needs to be resolved (ideally by comparing the WBS-5 survey methodology against actual RAC history) before the milestone goes to scrubbing. Alexei and Paolo also pressed a structural question: of roughly 18 new FY27 milestones, most read as R&D rather than operations, with little or no FTE attached, which contradicts Wei's own framing that the AFs barely have effort to cover account creation and emergencies — Rob's answer was that genuinely new integrations belong in 2.3.4 as a handoff point from 2.3.5/2.5 development into production, not as unfunded 2.3.4-native R&D, but that framing needs to be made explicit and consistent in the deck. Alexei closed the substantive discussion by naming the issue underneath all of this: WBS 2.3.4 (and ATLAS more broadly) still doesn't have an agreed answer for what the Analysis Facilities should look like for HL-LHC, and until that vision exists, individual milestones and FTE requests are hard to justify one at a time. Two smaller items landed as agreed fixes: quietly retiring milestones 537 (Condor glide-in federation, lost its champion when Lincoln left) and 538 (remote LGD access, no user interest per the WBS 5 survey) rather than carrying them as "delayed," and adding the two items Verena caught as missing — a milestone for installing the UC AF end-of-CA hardware, and a Tier-2-style slide showing what the hardware-aging picture looks like once end-of-CA purchases are folded in. The transcript excerpt ends as Wei raises a related, unresolved concern — that HPC/AmSC GPU access is creating an imbalance across AF users — without a proposed fix.

2.3.5 Continuous Integration and Operations (Ofer Rind, with Ivan Glushkov). Four L4 areas: 2.3.5.1 Operations Support, 2.3.5.2 Services DevOps, 2.3.5.3 Facility R&D, and 2.3.5.4 Facility Cybersecurity — new this year, implemented without dedicated funded effort in FY26 but with real preparatory work done. Total effort is 4.70 FTE in FY26 rising to a gross 5.95 FTE in FY27, of which 1.35 FTE is proposed to move to WBS 2.5, netting to 4.60 FTE remaining in 2.3.5. Accomplishments included Ivan Glushkov's appointment as ATLAS ADC Coordinator (adding US ATLAS representation in OAB, CREM, and WLCG Ops Coordination), completion of the Varnish migration for ATLAS conditions data (CREST/Frontier caching), the ADC analytics platform (Elasticsearch/Kibana) now running across the SSL River cluster at UChicago and CERN, and substantial agentic-AI development led by Ilija Vukotic — an OpenClaw/NemoClaw multi-agent harness under a new governance layer called "Shannon," a live HTCondor triage agent on the Analysis Facility, and a live AF AI assistant. On Facility R&D, RP1-core (the generalized, infrastructure-as-code successor to the original stretched-Kubernetes platform, developed this year by Aidan Rosberg) went operational, and the Open Data Facility (ODF) was deployed at UC as a from-scratch proof that the RP1-core pattern is portable to other sites rather than UC-specific — Rob gave the group's explanation of both. On cybersecurity, a real

four-part FY27 work plan now exists (Program Foundations, Site Security Coordination, Agentic Operations Security, Monitoring & IR Integration) addressing the gaps the 2025 Trusted CI assessment identified (no dedicated lead, no incident-response policy, no risk register, no stakeholder registry); a Federation Liaison Model stood up in May 2026 as a precursor to named security contacts at each Tier-2; and the team ran its first cybersecurity exercise with OSG Security and handled two live vulnerability advisories in June 2026.

The FTE tables needed the same kind of real-time reconciliation seen in 2.3.4. Shawn McKee's effort split between 2.3.5.1 and 2.3.6 (Management) had been mis-recorded last year (0.4 Management / 0.1 Ops rather than the agreed 0.3 / 0.2), an error that had already propagated into the master effort spreadsheet and needs a clarifying comment. Ilija Vukotic's apparent "increase" is not new funding — it's 0.2 FTE that was always his but had been sitting uncounted under 2.4.3 and never propagated to the spreadsheet; his real FY27 total (0.40 remaining in 2.3.5.2 plus 0.60 recolored to WBS 2.5, summing to 1.00) is flat versus FY26 once that 0.2 is accounted for. The person replacing the departed Lincoln Bryant on Facility R&D is Giordon Stark (UCSC) at 0.25 FTE plus another 0.25 recolored to 2.5 — Granola's transcript rendered his name as "Jordan," corrected here.

The most consequential finding of the session was on cybersecurity funding. The entire FY27 2.3.5.4 plan depends on hiring a 1.0 FTE cybersecurity lead under the new Cooperative Agreement, and the group initially reassured itself — with visible relief — that a full FTE had in fact been requested. Partway through the review, Ofer rechecked the actual NSF CA basis-of-estimate and found it allocates only 0.2 FTE for cybersecurity, matching Rob's existing effort rather than funding a new hire. This directly contradicts what the group believed and needs to be resolved before scrubbing, since if the 1.0 FTE isn't actually funded, the whole 2.3.5.4 FY27 work plan needs to be rescoped. Separately, Paolo found Rob's live explanation of how "Shannon" (privileged, infrastructure-facing agent governance) and "Elwood" (the user-space equivalent for AF users) relate — both developed and tested first on the disposable RP1 R&D platform before graduating to production — genuinely useful, but said the fact that it took 10–15 minutes to explain live is itself the problem: it needs to be pre-digested for a management audience, which Rob agreed to redo (an LLM-assisted rewrite was discussed, echoing the same fix already promised for 2.3.4). Paolo also pushed back that the 2.3.5.3 R&D activities read like "a collection of things we started working on" rather than being visibly anchored to ATLAS/WBS 2.4 priorities, and asked for explicit connections to existing milestones and requirements. On milestones, Ofer flagged that the 2.3.5.4 cybersecurity milestones aren't in the spreadsheet yet since there's no hire to assign them to; Paolo's guidance was that milestones generally shouldn't exist below about 0.2 FTE of labor and that the cybersecurity work plan should collapse to roughly one milestone per thrust (about four total) rather than one per activity — Rob committed to cutting the milestone count by roughly a factor of three across his areas. The group also identified a missing risk-register entry — no cybersecurity lead means no incident-response capability and possible forced site outages after a cyber incident — to be folded into the R8-13 risk Shawn had already proposed adding in the Overview session.

Closing discussion. After 2.3.5, Alexei gave a blunt overall assessment: WBS 2.3 cannot go into the real scrubbing with the 2.3.4 and 2.3.5 slides as they stand — he expects reactions

ranging from confused to hostile, and said neither talk yet delivers a clear message, describing them as collections of minor activities rather than a coherent case. His core ask: any FTE or milestone request has to be proportionate and legible to someone unfamiliar with the material — if the team needs 18 FTEs, ask for 18 with 18 milestones, not 1.2 FTE spread confusingly across many small ones — and both the slides and the underlying effort spreadsheets need real work before July 13, ideally within the coming week. He proposed an additional Level-3-only working session, tentatively Friday, July 10, between himself and Rob to go through the 2.3.4/2.3.5 material in detail. On scheduling, the group ran out of time for both 2.3.5 and 2.3.1 in one sitting; Tier-1 was moved to Wednesday, July 8 at 1:00 pm, replacing that day's regular facility coordination meeting, with Alexei presenting. The meeting closed on a recurring theme from 2.3.4 — why ATLAS needs Analysis Facilities at three separate sites at all — with Rob framing the three-site model as risk mitigation and flexible capacity-sharing (e.g., flocking idle Midwest Tier 2 cycles to the co-located Analysis Facility) rather than pure redundancy, while Alexei pressed that this framing needs to be backed by a real, demonstrable answer, since the question is close to certain to come up at the actual scrubbing. Rob noted he had been keeping his own Granola notes throughout and would fold them into the shared meeting-notes document.

Action Items

Action items are grouped by the WBS area they came up in, with cross-cutting items in a general list at the end.

WBS 2.3 Overview

Owner	Action	Context / Deadline
Shawn McKee	Check the Drive folder for pending sharing/access requests	Raised at start of Tier-2 section
Shawn McKee	Add risk R8-13 (facility cybersecurity incident, 2.3 as risk owner) to the formal risk register	Proposed in 2025 scrubbing, never entered; carry into updated register
Shawn McKee	Add new risk covering aging/robustness of facility hardware given the "run it until it breaks" mitigation strategy	Suggested by Paolo during Overview risk discussion
Alexei Klimentov	Clarify in 2.3.1 slides why some 2.3.1/2.3.4 effort numbers differ from the master effort spreadsheet (effort from non-2.3 people paid from 2.3 sub-WBS)	Raised during Overview FTE summary discussion

WBS 2.3.2 — Tier-2 Centers

Owner	Action	Context / Deadline
Fred Luehring	Recheck with Jono whether a no-cost extension is possible for FY26 equipment funds; drop the bullet from the slide if confirmed unavailable	Encumbrance deadline is November 2, 2026; Fred has already been told "no" once
Fred Luehring / 2.3.2 team	Add explicit note on the FY27 hardware capacity slide that pledged targets already assume equipment purchased with end-of-CA funding (i.e., not optional)	Requested during Tier-2 hardware review
Fred Luehring / 2.3.2 team	Consider forming a small task force to evaluate storage purchasing and warranty strategy (e.g., pre-paying for 7-year warranties) in light of a possible one-year Long Shutdown 3 extension	Raised during Tier-2 FY27 plans discussion

WBS 2.3.3 — Heterogeneous Integration & Operations

Owner	Action	Context / Deadline
Rui Wang	Revisit milestone 2.4.2-614 (PanDA/Triton inference service) before finalizing "retired" status; confirm whether similar capability already exists at another site and reconcile with related milestone 2.4.1-18	Flagged as needing a follow-up note during 2.3.3 milestone review
Rui Wang	Update FY27 milestone language to reference the IRI API using correct DOE terminology rather than generic "Globus Compute APIs"; confirm exact wording with the reviewer via comment	Raised during 2.3.3 FY27 priorities discussion
Alexei Klimentov, Rui Wang, Doug Benjamin	Take the Globus Compute vs. IRI API / HPC integration strategy discussion offline	Alexei stated he was not satisfied with the in-meeting answers
Doug Benjamin, Rui Wang	If considering a broader Globus overlay-cluster approach or other FTE-scope changes for 2.3.3, secure explicit buy-in from Zach and ADC/project leadership first	Raised in response to Paolo's question about whether more FTE would be used differently

WBS 2.3.4 — Analysis Facilities

Owner	Action	Context / Deadline
Rob Gardner	Add a milestone covering installation of the UC AF end-of-CA hardware	Verena caught this as missing from the 2.3.4 deck; confirmed as an omission
Rob Gardner	Add a Tier-2-style slide/table showing the UC AF hardware-aging and warranty picture once end-of-CA purchases are folded in	Verena specifically referenced Rob's aging-hardware slide from elsewhere and asked for the equivalent post-CA-refresh view; Rob agreed both this and the item above were missed
Rob Gardner / Wei Yang	Delete the GPU-shortage complaint bullet from the UC AF FY26 Challenges slide before the real scrubbing	Rob asked for this cut live in the meeting to avoid inviting unnecessary reviewer questions
Rob Gardner	Resolve the Local Group Disk contradiction: Wei's slide (based on a WBS 5 user survey) treats LGD as underutilized and proposes reorganizing it, but Rob has attended ~20 RAC meetings over 2+ years without a single space complaint	Needs reconciliation (e.g., comparing WBS 5 survey methodology against RAC history) before the LGD milestones go to the real scrubbing
Rob Gardner	Rebalance/clarify the Chicago (Hu) FTE line in the 2.3.4 effort table — consider moving BinderHub/federated-K8s/notebooks/Dask/HTCondor and ServiceX items out, since ServiceX is already gated by and counted under 2.3.5.2 and the rest leans on Midwest Tier 2 support	Paolo flagged the Chicago portfolio as disproportionately large/scattered next to BNL AF and SLAC AF; Rob agreed in the meeting
Wei Yang	Get Qiulan Huang to debug the BNL AF monitoring database — BNL's accomplishment-slide numbers (SSH/Jupyter/batch users, jobs, hours) look inconsistent with the other AFs and changed between the June 11 snapshot and a re-pull the day before the meeting	Already requested; awaiting response as of this meeting

Wei Yang / BNL	Prepare utilization plots/tables showing that BNL AF and SLAC AF's smaller, resource-constrained footprints are actually saturated, to have on hand if a reviewer asks why constrained sites keep getting funded despite falling user counts	Raised by Paolo during accomplishments review
Wei Yang / Shawn McKee	Reconcile the 2.3.4 FTE table against the master effort spreadsheet: fix the FY26→FY27 total (1.05→1.25 doesn't include Pelosi's new 0.1 FTE, which per Alexei's math should make it 1.35); add Nakata to the FY26 effort spreadsheet (present on the slide but missing from the spreadsheet); update the name of the person Nakata replaced in Chuck's spreadsheet	Discrepancies identified live during the FTE summary slide
Wei Yang	Confirm BNL's new 0.1 FTE for Pelosi with BNL leadership — Alexei said this addition was never discussed with BNL	Raised by Alexei during the FTE summary
Wei Yang	Do the hardware-request "homework": show current hardware is saturated, show its age profile, then justify a multi-year (FY27/28/29) purchase plan, rather than presenting a flat replacement-cost figure	Paolo declined to accept the current framing of the ~\$331K UC AF (plus BNL/SLAC) FY27 hardware ask as-is
Wei Yang	Investigate and resolve apparent duplicate milestone numbers 539–542 in the milestone spreadsheet	Shawn McKee flagged these as duplicating pre-existing entries of unclear origin
Wei Yang	Convert 2.3.4 milestones that use "ongoing" or "quarterly" as the date field into fixed-period, reissue-next-period milestones (matching the Tier-2 pattern)	Shawn McKee flagged inconsistency with how recurring milestones are tracked elsewhere
Wei Yang	Standardize milestone/effort table formatting across the 2.3.4 deck	Verena asked for consistency (e.g., matching the FY27 "MS ID / Description / Effort / Baseline / Impact" table style used elsewhere in the deck)

Wei Yang	Add explicit FTE estimates and a clearer ops-vs-R&D framing for the FY27 "big picture" items (milestones 2.3.4-15 through -18: hardware refresh, storage quota/policy, LGD reorganization, HPC/AmSC GPU utilization, general-purpose AI inference)	Alexei and Paolo both flagged that ~18 new milestones read as R&D with little/no attached effort, in tension with the claim that 2.3.4 has no spare effort beyond core operations
Wei Yang	Retire milestones 537 (Condor glide-in cross-AF federation — lost its champion when Lincoln departed) and 538 (remote LGD access — no user interest per WBS 5 survey) cleanly, with a note that a revised FY27 proposal may revisit the same goal using the IRI/AmSC interface instead	Agreed live between Rob and Wei
Wei Yang	Decide on and justify an approach for general-purpose AI/LLM access for AF users (commercial subscription at ~\$20–250/user/year vs. self-hosted inference vs. NRP-style open-model access)	Wei flagged this as unsolved and high-risk; Rob suggested evaluating NRP (already integrated into Jupyter) alongside SLAC's "I ² " and a Berkeley offering (currently staff-only) as alternatives to a blanket commercial subscription
Wei Yang	Add a note to the HPC/AmSC batch-GPU milestone requesting collaboration with 2.3.3 (Rui Wang / Doug Benjamin), given their report that NERSC GPU queues are largely underutilized	Raised by Rob as a possible quick win/backfill source

WBS 2.3.5 — Continuous Integration and Operations

Owner	Action	Context / Deadline
Ofer Rind, Rob Gardner	Urgently re-verify the NSF CA basis-of-estimate for the cybersecurity hire — confirm whether it actually funds 1.0 FTE or only the existing 0.2 FTE from Rob — before presenting the 2.3.5.4 FY27 plan at scrubbing	Ofer found a live contradiction mid-meeting between the group's belief (a full FTE was requested) and what the CA proposal document appears to show (0.2 FTE only); flagged as possibly discussable Wednesday

Ofer Rind / Shawn McKee	Add a clarifying parenthetical note reconciling Shawn McKee's 2.3.5.1 / 2.3.6 (Management) effort split — should be 0.3 (Management) / 0.2 (Ops) for FY26, not what shipped last year	Error had already propagated into the master effort spreadsheet
Ofer Rind	Note in Ilija Vukotic's FTE line that the FY27 total is flat once accounting for 0.2 FTE that was always his but sat uncouncted under 2.4.3; get this propagated to the master effort spreadsheet	Confirmed against the spreadsheet live in the meeting — spreadsheet still shows the old 0.8/0.2 split
Ofer Rind	Confirm Giordon Stark's (UCSC) 0.25/0.25 Facility R&D / WBS 2.5 split as Lincoln Bryant's replacement, and correct the name in any notes (Granola's transcript rendered it as "Jordan")	Identified during 2.3.5.3 FTE table review
Rob Gardner	Redo the Shannon/Elwood/RP1 explanation slides for a management audience (same note already given on 2.3.4); consider running draft language through an LLM for plain-language translation	Paolo found the live 10–15 minute explanation itself to be the problem, even though the content was useful
Ofer Rind, Rob Gardner	Connect 2.3.5.3 Facility R&D activities explicitly to ATLAS/WBS 2.4 priorities and existing milestones/requirements rather than presenting them as a standalone activity list	Paolo described the current framing as feeling like "a collection of things we started working on"
Ofer Rind	Consolidate 2.3.5.4 cybersecurity milestones to roughly one per work-plan thrust (~4 total, each representing at least ~0.2 FTE of labor) instead of one per activity; add priorities to the milestone spreadsheet	Paolo's general guidance on milestone granularity; Rob committed to cutting milestone count by roughly a factor of 3 across his areas
Ofer Rind, Rob Gardner	Add a formal risk-register entry for "no cybersecurity lead → no incident-response capability / possible forced site outages," folding into the R8-13 risk Shawn proposed in the Overview	Gap identified live during the 2.3.5 risk review

Ofer Rind	Consolidate the dense, jargon-heavy 2.3.5.2/2.3.5.3 milestone/table slides generally, per Paolo's feedback that some boxes are too dense to absorb while listening to a talk	Raised during cybersecurity milestone review
-----------	--	--

General / Cross-Cutting Action Items

Owner	Action	Context / Deadline
All WBS 2.3 L3 areas	Submit FY27 budget requests as a full budget under 2.3 (not pre-split for WBS 2.5); flag in comments what should move to 2.5	WBS 2.5 currently has zero allocation; coordinated move to follow later, per guidance given to all Level 2 areas
Rob Gardner, Shawn McKee, Alexei Klimentov	Address the standing strategic question of what the Analysis Facilities should look like for HL-LHC before individual 2.3.4/2.3.5 milestones and FTE requests can be fully justified	Alexei named this as an unresolved, multi-year-old fundamental question underlying the milestone-by-milestone disputes in this session
All L3 presenters, especially Wei Yang and Ofer Rind	Rewrite dense, jargon-heavy slides for a management-level audience, and rebalance milestone/FTE requests so they are proportionate and legible to someone unfamiliar with the material, before the July 13 scrubbing	Alexei's blunt closing assessment: 2.3.4 and 2.3.5 as presented would likely draw a hostile reaction; asks should be sized to match ("if you need 18 FTEs, ask for 18," not fragments spread across many small lines)
Alexei Klimentov	Present WBS 2.3.1 (Tier-1) on Wednesday, July 8, 2026 at 1:00 pm, replacing that day's regular facility coordination meeting	Agreed at the close of this session; Alexei needs to create the Indico entry
Rob Gardner, Alexei Klimentov	Hold a dedicated working session to rework the 2.3.4/2.3.5 slides together	Tentatively proposed for Friday, July 10, 2026, ahead of the July 13 scrubbing; participants should keep that day clear
Rob Gardner	Post Granola notes from this session into the shared meeting-notes document	Committed at the close of the meeting

Transcript

Speaker labels below combine the hints Rob provided with content-based inference, since the Granola transcript included no speaker diarization. Names are used where the content anchors them with reasonable confidence; "Unidentified speaker" is used elsewhere, with a best guess noted in brackets where the context strongly suggests one of "Paolo" or "Verena."

Administrative Opening

Shawn McKee: Let's go ahead and get started — a few more people still need to join, but Wei said he'd be late. Everyone should have seen today's Indico page. At the top there are links to the Google folder with all the meeting information, the WBS 2.3 milestone spreadsheet we've been updating, the Indico page for next week's scrubbing, and a shared meeting notes document that everyone should have edit access to — please help note any action items or important changes coming out of each presentation. Today's agenda shifted slightly: Alexei has to leave early, so we moved Tier 1 to the end. He also asked that if he's not back by 1:00 when 2.3.4 is scheduled, we swap 2.3.5 in ahead of it. Anything else for the top of the meeting? All right, let me share the 2.3 overview slides.

WBS 2.3 Overview (Shawn McKee)

Shawn McKee: What I'd like to do isn't so much present the slides as walk through them and get comments and discussion — what's missing, what needs to change, or any other action items we need to take. This first presentation gives an overview of all of 2.3. Up front, I want to note these slides are meant as a quick overview to set the stage; I'd like to keep the detailed questions for the individual Level 3 presentations later. So there's the "apple pie and motherhood" statement at the front: what do we do as a facility? We deliver resources, we provide the US ATLAS physics community with additional resources, and we responsibly engage in community R&D to improve things — especially with HL-LHC in mind.

Unidentified speaker [Paolo]: Shawn, feel free to fly through the slides — except when we get to the FTE numbers. Don't feel like we have to go through the whole talk.

Shawn McKee: Just a quick slide about the unusual, challenging circumstances we're experiencing. Comments are shown here — if people think anything else should be called out, let me know.

Unidentified speaker [Verena or Paolo]: These are, of course — "challenging circumstances" may be selling this a bit short. These should be highlighted. Is there anything in what follows that reflects these huge cost increases, and are we assuming for now that things will stay as they are?

Shawn McKee: There are certainly things in the Tier 2 and Tier 1 slides where we talk about that, and I'll mention it later — but yes, that is going to be our big challenge. Let's move on. This is the updated org chart since last year — I've marked all the changes in red. There was one

change to 2.3.5.3 just today. I also noted that there's an ATLAS-wide position, Ivan's position, which I've circled in blue.

This is probably the first slide that will generate some discussion — the overall FTE summary for FY26 and FY27. I pulled these numbers from the effort spreadsheet and from each presentation; the comments show which slide in each presentation the numbers came from. I noted that some of the 2.3.4 numbers don't match what's in the actual effort spreadsheet, and there's a similar discrepancy for 2.3.1 — the spreadsheet has different numbers than what's shown here. Alexei, you had a comment, but our goal was to make sure the spreadsheet is the source of truth and the slide reflects that. Go ahead, Alexei.

Alexei Klimentov: I can address it when I present — I asked [Shawn] explicitly about this. We had effort from people who are not formally in 2.3, but we paid for that effort and are counting it. I believe it's fair to count it because they contribute to 2.3, but I'll address the details in my slides.

Shawn McKee: Okay — I'll add a brief note on 2.3.1, because I know what happens when I show this slide: people check it against the source spreadsheet and ask why the numbers don't match. On the slide I've color-coded things — green where effort increased, red where it decreased, and purple for assumed movement of effort to WBS 2.5, which hasn't been finalized yet. Those are the highlighted movements. All of 2.3 — just to clarify, this is a question we were asking every single Level 2 area: if WBS 2.5 did not exist, would you have 27.12 FTE in FY27, or 29.17?

Alexei Klimentov: We would have 29.17.

Unidentified speaker [Paolo]: Okay, that's not going to happen — but we can certainly ask for it. Just be aware, this isn't the climate to ask for a 10% increase going from 27 to 29.

Alexei Klimentov: One of those FTEs was the new cybersecurity position.

Unidentified speaker [Paolo]: Right — just to avoid any confusion, what we've suggested to all Level 2 areas, and the same applies to 2.3: submit a full FY27 budget, because WBS 2.5 right now has a zero allocation. Then put in the comments what you'd like to see move to 2.5 out of the FY27 request, and we'll come back and discuss moving particular pieces with you. We know Ops is going to move for sure, but even for HIOPS, I'd say leave it in FY27 under 2.3, and we'll tell you when we're ready to actually make the move — otherwise it gets very confusing. That's what everyone else across 2.4 and 2.3.5 is doing, so let's do that consistently.

Shawn McKee: Sounds good. I'm putting a comment so we don't forget the 29.17.

Moving on — accomplishments. Five statements here about some important things. This is delivered Tier 1/Tier 2 CPU resources; I've noted where US ATLAS stands. The left panel is all completed jobs, the right is all successful jobs, EPSCoR-weighted — you can see the percentages.

Unidentified speaker [Paolo or Verena]: Sorry to bring you back — the facility accomplishments are written in nice English, so they're easy to understand, but I was wondering: "facilities adapting to opportunities from the rise of LLMs" — can you be more concrete about what that means?

Shawn McKee: There's a lot of work incorporating LLMs into operations, planning, and user support — that's been an ongoing activity.

Unidentified speaker: I see — so the opportunities are in terms of operations, not opportunities in terms of serving LLMs?

Shawn McKee: Right, I should make that clear — it's us taking advantage of LLMs for our own efforts, not serving LLM workloads.

Unidentified speaker: Thank you — I'll leave that as a comment so you can find it later.

Shawn McKee: Let's come back to this slide — on the left, US delivered 22% of all completed jobs; on the right, for successful jobs, we delivered 26%. That's just what the numbers show. Here are the analysis jobs completed — we look at these in particular because they're difficult, and you can see where the various sites fall this year; we account for 25% of analysis jobs. This is storage resources delivered — the US is right around 25%. We had a number of challenges: availability and cost issues for needed storage and memory, the cost increases we're seeing, a lot of work and discussion around Analysis Facility convergence, the EL7-to-EL9 transition, all the DC27 prep work and feedback to the facilities, and the cybersecurity gap. There are also continuing issues with personnel and expertise — difficulties and delays filling positions, though we've had some successes. It takes time and effort, and requires skill sets that are difficult to find at the salaries we can pay.

Here's a summary of all the milestones. Note: there are some new 2.3.5 milestones not yet in this table because they need baseline and end dates before my script can pick them up. Modulo that, this should be all milestones relevant since FY25 — basically since last July 2025 — including new ones listed as proposed. There are 187 total, and they break down this way.

Unidentified speaker [Alexei]: Shawn, have you included the milestones for the effort proposed to move to 2.5, or not?

Shawn McKee: Yes — I believe we've put even things we think will move to 2.5 into the 2.3.x milestone list.

Unidentified speaker: Okay, so the only changes needed are to the effort numbers.

Shawn McKee: Right. Okay, good — hopefully that's not too bad. Thanks — we did do a bit of work here, and when we get to the Level 3 presentations we'll show which milestones we think would move. Okay — high-level plans. We're moving into LS3, so there are opportunities and challenges here. At the top is delivering our pledges and resources, which is challenging — it will require creative solutions and some hard choices. Related to that is preparing and executing

DC27, the cybersecurity issues, migrating suitable effort into the new WBS 2.5 area, continuing our collaboration, and continuing to build our bench. These are somewhat standard items, but important to call out.

Risk analysis — here are the 2.3 risk updates. I've added a sentence or two under each describing current relevance. One thing that still needs to be done: these need to be translated into actual changes to the risk register — risks that increased, mitigations that changed, or probabilities that shifted all need to be captured and handed off. We proposed adding risk R8-13 last year and it never got entered — I still think it's an important risk: that the facility experiences a cybersecurity incident, with 2.3 as the risk owner. We discussed last year where it would go. As before, I've got a document where new risks from any L3 area can be discussed, fleshed out, and agreed on. Right now R8-13 is the only one in there.

Unidentified speaker [Paolo]: I'm curious about the hardware cost risk, which is clearly realized — what were the mitigations, and were they triggered?

Shawn McKee: One second... The mitigation was either spending more money or delivering fewer resources. One creative option is keeping equipment longer — that's the default one we're pursuing: not retiring equipment just because its warranty is up, as long as it's still functioning. It's risky, but what else can we do?

Paolo: Okay — there are three new milestones, and now you have a new risk about hardware, essentially — the robustness of our equipment is a risk, especially as it gets creakier and creakier.

Shawn McKee: Thank you, Paolo — I've been meaning to suggest a risk exactly like that and kept forgetting. That's a good idea.

Unidentified speaker [likely Kaushik De, SWT2]: In terms of actually mitigating the risk, at Southwest Tier 2 we're extending warranties — it costs a bit more, but it compensates for the higher quotes for a while.

Shawn McKee: Right, and there are other variations too — if you have old retired equipment, you can keep it as spare parts, if you've got the right models on hand.

Paolo: What about using more HPC resources as backfill for pledge, instead of just beyond-pledge? That's another mitigation option.

Shawn McKee: Well, then we'd need it defined as pledged, and that has a requirement — we'd need guaranteed use of that resource to actually pledge it. Let's move on. So, three new milestones were added: the facility completes participation in DC27; the facility completes post-DC27 work — analyzing the results and summarizing the implications for the facility, which will take a while, so I put that about six months after DC27 wraps up; and the third is deploying all our pledged resources, including the end-of-CA purchases, which I put in April.

Unidentified speaker [Verena or Paolo]: Is it worth having the end-of-CA purchases as a separate milestone?

Shawn McKee: We could — I thought about that. It could end up as a separate action depending on when we get approval; for example, if we need a no-cost extension, it could happen later than planned.

Unidentified speaker: Yeah, but for now that's the plan. Also, why is the priority "medium" for deploying the CA resources?

Shawn McKee: At that point we're not below any pledges — we're still okay, so this one can shift. That was my justification.

Unidentified speaker [Paolo]: Well, there's also the directive that we don't want too many milestones marked "high" — but come on, that seems significant; if we miss the end-of-CA purchase, we'd really be in trouble.

Shawn McKee: I know — but this milestone is about deployment, not acquisition. It's specifically about getting it deployed by this date, which is why I called it medium — six months later we're still fine, okay. Question for you, Paolo — you pointed me at a particular talk and asked for a specific format for how milestones should be presented. This isn't in that format — should these be in the format you recommended?

Paolo: I need the information more than the format. I'd pointed you to a talk — I think in 2.4 — about priority on a one-to-five scale.

Shawn McKee: Right — separately, I normalized all the milestones for the milestone summary pie charts: I mapped priorities 1 and 2 to "High," 4 to "Medium," and 5 to "Low."

Paolo: Right, I didn't say "US ATLAS High" or "ATLAS High" specifically.

Shawn McKee: Right, you didn't — so that's the translation I used.

Paolo: I'd say this format is okay — the reason I needed that specific format was to make sure I could extract milestones from the slides with a prompt reliably. I think this will work, but I may ask you to adjust something.

Shawn McKee: I've got all my new milestones in the format you wanted — that saves me work and eventually may save you work too. I was also going to go back and put the ones we completed this year in that format.

Paolo: That's not necessary — the completed ones don't matter for the format.

Shawn McKee: Okay, this is the last slide. To summarize: we delivered; with LS3 we have opportunities; we need to think about how risk-averse we are regarding cybersecurity; and now we'll move into the Level 3 presentations. That's it.

Unidentified speaker [Paolo or Verena]: Sorry to be a pest, but whoever owns the Drive area — maybe that's Shawn — can you check for sharing requests? That would help.

Shawn McKee: Yep, I'll do that. Meanwhile, let's go on — since we're jumping over Tier 1, next is the Tier 2s.

WBS 2.3.2 — US Tier 2 Centers (Fred Luehring)

Fred Luehring: If you don't mind, I'll share my own screen — that way I can jump around. And Rafael wasn't able to join, right? I haven't heard anything from him; he's stuck in transit in India. Let me get the right window here — okay, can everybody see that? Before I go full screen... right, let me remember — you should have access; someone else got there before I could.

Unidentified speaker: Slideshow, top right, where it says "Slideshow."

Fred Luehring: Thank you — that was too easy. Okay, so this is the Tier 2 centers. They're less software and more hardware than many of the other areas. Personnel is almost unchanged. Midwest Tier 2 shifted—

Unidentified speaker: Sorry, I don't see the slides on Indico — are they in the folder?

Shawn McKee: Yeah, they're in the folder — let me get to my phone, give me a second, I'll link them there too. Sorry, I forgot to do that. I've pasted the link there as well, as you asked.

Unidentified speaker: Yeah, I saw you'd jumped in — we're already looking at them, thanks.

Fred Luehring: Okay, I'm on slide 2, personnel. We moved Fengping [Hu] off of Midwest Tier 2 — that effort went to the Analysis Facility — and brought on 0.2 FTE of Judith [Stephen] to make up for that move. So the upshot is we were at 9.62 FTEs across the entire complex this year, and we expect to be the same next year.

Unidentified speaker: So the percentages weren't 100% for some individuals?

Fred Luehring: That's right — they're in other projects. I only commented where I knew they were in 2.3 — basically, they were around 2.3 somewhere. This slide is a little out of date — I made it before June 30th. Before we present [to the review], I'll show the entire year; the numbers will change very slightly.

Unidentified speaker: So I guess that means FY26, right — not FY25? And the other question — sorry — Shawn already showed this, right?

Fred Luehring: I showed the whole facility, which includes Tier 1.

Unidentified speaker: Oh, I see — that's the difference.

Fred Luehring: Okay, and these up here — this is something Rafael introduced last year — are links to the underlying plots so you can check them. Next. Now we get to the slide where I always end up wrangling with the reviewers at length. It's particularly complicated this year because of how the funding has gone — we've basically not spent much of our FY25 funding, and now we're going to get a further infusion for FY26, on top of that. The upshot is there's a new infusion of \$1.7 million per site for equipment. Currently we're in good shape — there's no question we're covering the US target and then some, in all cases, for both compute and storage. I'd been using "greater than 7 years" for old equipment, but since Rafael uses "greater than 8 years," I switched compute to greater-than-8-years for consistency — you can see the amount of really old equipment this year is very small. We generally haven't had problems — we lose a few percent of compute that's this old, it's not dying en masse. The real drawback is, even though we don't pay for electricity on it, it's an expensive way to get CPU cycles. Storage is in okay shape too — because we've put so much effort into storage purchases, we're not nearly as behind; we have a much better trajectory on older storage. In fact, nobody is running storage older than six years and counting it toward the target.

Okay — before I leave this slide, I want to note: this was the last year that Midwest Tier 2 was 1.5 times bigger than the other Tier 2s, so that creates some funkiness in the FY27 pledges, because now the pledges and targets become equal across sites.

Unidentified speaker: Is this what happens when the targets and pledges have to be equal? Targets are what the US gets, and pledges are...?

Fred Luehring: Sorry, let me be clear — both sides, the targets, are the same now, not the pledges. Midwest Tier 2 is no longer 1.5x different. So effectively, the other three Tier 2s get a significant increase in pledge, and Midwest Tier 2 — which was already sitting well above the 1.5x — looks exceptionally, exceptionally rich. That's why I wanted this up here; I can just imagine having to discuss that with the reviewers, and I wanted to be sure I made it clear. For computing, I'm assuming we don't retire anything in FY27 — we just run it until the wheels fall off — and we don't buy any new compute in FY27. So we end up more or less meeting target everywhere.

Unidentified speaker [Paolo or Verena]: Does this include the pledged hardware we may buy with the end-of-CA funding?

Fred Luehring: Yes, of course.

Unidentified speaker: Okay, good — maybe note that somewhere so it's clear, since we need to emphasize that it's not optional.

Fred Luehring: Okay, yeah — I have detailed backup slides showing all the calculations, but I'll add some commentary here explaining what's going on. You'll notice I mark whether it's FY27 or FY26 — when I get into FY27, everything is against the FY27 targets; the previous slide was against FY26 targets. I see I have a typo there, despite spending hours on this. We don't know exactly what we'll get, but we're assuming a budget of about \$180,000 per Tier 2 for equipment.

Following Kaushik's instructions, to keep things simple I took two-thirds of that, since only two-thirds of the new CA funding falls in FY27 and the money is allocated per calendar year — that brings it down to \$120,000. I'm sticking with \$100/terabyte for storage. We should move on — we're pressed for time. Sorry.

FY26 has been a challenging year for two reasons. First, the funding level for equipment purchases was unknown until mid-May, when it was set at \$1.7 million per site. NSF has formally approved this, and the paperwork is in place so we're not restricted by the budgets we submitted five years ago. But NSF has not yet actually permitted us to spend the money. So we have to encumber the funds — that is, issue POs — under the current rules, by November 2nd. And of course we're dealing with the risk already discussed: significant price increases, in some cases four to six times, for certain types of memory.

As an action item, we should double-check whether we can actually get a no-cost extension — I believe the answer is no, but Amir suggested it might be possible. I wouldn't have put it on the slide otherwise.

Unidentified speaker: That's why I'm saying, let's check again offline with Jono.

Fred Luehring: I did check, and he told me no — but we'll check again, and if not, we'll drop this bullet.

Unidentified speaker: It's written so it can easily be dropped. It's probably better to drop it, since there are also risks associated with requesting it — put that in the minutes.

Fred Luehring: Yeah — here I show what the hardware purchases are planned to be; most sites are buying a lot of storage, because that's the hard thing to bring up quickly when demand increases, and it ages out faster. We have a lot of ongoing M&O activities, and the remaining 25% of admin time is spent on non-routine things. For the 2026 milestones, we did extremely well — we closed a bunch of long-standing ones, and were delayed on only three minor ones. The backup slides include the descriptions. For the new milestones, this shows by site how many there are — 14 of the 23 new milestones are for M&O activities; we're trying to do a better job of reflecting M&O. I have these milestones in the backup slides in the requested format, though I added one additional line for priority and status.

Unidentified speaker [Paolo]: Good idea — actually we should have done that across the board, since I've specifically been asking about how many are M&O.

Fred Luehring: Yeah, we should have — it's too late for this round, but you'll have to change the script that reads the milestones.

Paolo: How did the effort assignment by milestone work over the last year?

Fred Luehring: The way we're set up, we have the right amount of effort at each site — well, people work hard, let's put it that way. A lot of it happens in small increments. There were some delays caused by staffing shortfalls, but we've managed to eventually cover every milestone that

was delayed — sometimes we just have to push lower-priority things aside for higher-priority ones. But in general, the message of this slide is: all good. Let's move to plans. We have the right people, and we'll complete the milestone items along with our normal day-to-day activities. We're not planning any major new activities beyond what we normally do, and the milestones reflect that. The first part of FY27 will be quite busy procuring and installing about \$2 million worth of equipment per site. There will also be some effort during that period on preparing for the data challenges — and since several of the large procurements are network-related, that timing fits in well.

Paolo: Reminder — we're not reading the slides verbatim, so let's go through the points you think could be trouble spots.

Fred Luehring: Moving on — there are two existing risks plus the one you suggested. I don't have anything to add beyond that. This just shows explicitly the most recent memory price plots — if you're buying fifth-generation AMD CPUs, the current generation, you're buying one of these two types of memory. There's a huge range in prices in this survey, but the average is more or less flat for 16GB and gradually rising — though there was a period where it shot up like a rocket. I haven't drawn any conclusions from it.

Unidentified speaker [Paolo]: Just a question — should we have a summary/conclusion slide for each L3 area? Usually something like: here's the end, here's what I want you to take home.

Fred Luehring: The way the talk is structured, the message — the takeaway — is what we're asking for, because of the hardware cost. Can I write that on the last slide? Unfortunately it's not really a laughing matter, but I have to laugh or else I'd cry — because we finally get a really nice infusion of money, and it feels like the stars are aligned against us. Okay, that's it.

Unidentified speaker [Paolo]: My overall impression from 2.3.2 is that if it weren't for this cost crisis, you'd be doing swimmingly.

Fred Luehring: Basically, yes — if the budget's there, we're delivering everything, and the amount of very old hardware isn't increasing significantly. This plot is probably the main negative point. One other thing: because five years lines up with the start of Run 3, we're about to start losing a lot of storage to aging out — but we've covered ourselves on that by pre-buying a bunch of storage ahead of time. Several sites are going to lose five or more petabytes of old storage if we don't keep running it past warranty. I should probably put that on the final slide too — since we're only looking at FY27 here, you don't see it, but I'll get the number of how much goes out of warranty in FY28.

Unidentified speaker [Verena or Paolo]: A random thought — it's very likely that Long Shutdown 3 will be extended by one year. What would the consequences of that be for the Tier 2s? Negative?

Fred Luehring: It would largely be positive, I think, because we wouldn't need to ramp up as fast, even though we're not necessarily getting more money. It is true, though, that an extra year means some fraction of the storage we bought goes out of warranty and gets older than six

years. What we can do — if it's bought from Dell, for example — is pay a significant amount to extend the warranty. That's probably the middle ground. That's the card we've got in our pocket — right now most sites have storage on five-year warranties; I think Kaushik said he's moving to six. Our future purchases were planned on seven-year warranties, I think, for almost everybody.

Unidentified speaker: Right — and if the big storage purchases we're making this year with the extra money are on seven-year warranties, that would handle a one-year extension.

Fred Luehring: That's probably an offline action — we should think through all the implications, that'd be good to understand.

Unidentified speaker [Paolo]: We've talked about putting together a small task force to think about this purchasing conundrum.

Shawn McKee: I think that's a wise thing to consider doing. Okay, let's move on to 2.3.3 — we're almost on time.

Fred Luehring: I'm going to stop sharing here.

WBS 2.3.3 — Heterogeneous Integration & Operations (Rui Wang / Doug Benjamin)

Shawn McKee: All right, on to 2.3.3 — are you able to share?

Rui Wang: Yeah, just give me a second... I'm showing the correct screen. Okay, can I go full screen?

Unidentified speaker: Yeah, looks good.

Rui Wang: Okay, for 2.3.3 — should I go through the accomplishments?

Shawn McKee: Only if you think there's something exceptional to call out.

Rui Wang: Okay, so this combines 2.4.2 and 2.3.3. The main items are automated workload management using Globus Compute with the NERSC Superfacility API and Harvester, in collaboration with HEP-CCE, and the Rucio-plus-Globus test, which showed success in collaboration with HEP-CCE — those are the two main items. Let's move to the FTE summary — it's pretty simple for 2.3.3: in FY26 the FTEs are combined between 2.3.3 and 2.4.2, giving a total of 0.7, and in FY27 that number stays the same. The current proposal is that they move to WBS 2.5.

Shawn McKee: I see you're marking your other areas, but not for Doug — let me see... yeah, 0.6 non-ATLAS, that's easy to note. What are the other areas for you, Doug?

Doug Benjamin: 0.6 non-ATLAS for my remaining effort, and the other piece is 2.4.1.

Shawn McKee: Ah, that's 2.4.1 — easy.

Rui Wang: Okay, go ahead — challenges for FY26, mainly around workload efficiency: empty pilots, which is the highest-impact issue right now — let me go through the description. This may actually be solved by ADC giving us more tasks, so we just need to make sure we keep getting tasks — that might move this from a challenge to an accomplishment.

Unidentified speaker (joking): You've got one week to figure that one out.

Rui Wang: Doug, let me know if you think this should move to "countermeasure" or stay listed here as a challenge.

Doug Benjamin: I'll work with you to word it right.

Rui Wang: We also have "edge job failures" — pilots close when the SLURM allocation ends. We're going quickly through the slides, we don't need all the detail. The other thing: we need to reconnect Globus to BNL's production dCache so we can transfer data again.

Doug Benjamin: Otherwise we can't transfer data — for TACC, and for ALCF as well; if we want this working for FY27, we can't transfer any of it without the reconnection.

Alexei Klimentov: Sorry, I may be misunderstanding something — how are we running now? Are we using Rucio at all?

Doug Benjamin: TACC isn't running right now, so there's no transfer that way. We're okay — we're using the pilots to copy data directly from BNL's dCache and put it back. NERSC hasn't caught us on this yet.

Alexei Klimentov: But it's not what they recommend.

Doug Benjamin: Okay — the recommended way would be to use Globus. Fine.

Rui Wang: Now, FY26 milestones — it's busy here. We have the two allocation milestones completed. A few were retired because Frontera was retired, so those milestones were retired, and the others are all on schedule. The first milestone — hyperparameter optimization — "no existing use case." That one was delayed because there's currently no use case; once there is one, it can be reopened.

Unidentified speaker [Paolo]: "Delayed" sounds bad — it's not really delayed if there's no real use case. Are you saying we should retire it instead?

Rui Wang: There's no reason to look bad — let's just retire it, that's what I'm thinking.

Paolo: "Delayed" sounds like you weren't able to do it, which isn't the case — but I don't think we have a "paused" status; retire it, and we can reactivate or create a new milestone if the need arises.

Rui Wang: Okay — for the [2.3.3.2] milestones, the first is retired because we don't see a clear strategy for how to do it. There's a related new milestone in 2.4.1, so maybe this becomes a joint milestone with 2.5 in FY27. The other two are on schedule but linked to the Globus connection challenge.

Unidentified speaker [Verena]: Just out of curiosity — why are these all low priorities? I'd have thought getting these PanDA queues connected was of interest.

Doug Benjamin: They're low priority because these are HPCs — by the old definitions, HPCs are opportunistic. If ALCF or OLCF don't happen, there's no negative implication for our daily deliverables. If we stopped using NERSC, Jeremy would be annoyed, and we'd lose about 3% of our computing — so it's not just that Jeremy would be annoyed, but that too.

Unidentified speaker [Verena]: So the point is these would be nice-to-have activities — we may need them given the cost discussion earlier in this meeting, but right now they're optional. And the first one, about using PanDA to operate an inference service — you retired that one?

Doug Benjamin: Yes, because it was proposed too early, technically — we were probably a year and a half too early.

Unidentified speaker [Paolo]: Okay, so this might be one that's retired now but could come back in the future — similar to a comment on a similarly phrased milestone elsewhere in 2.4.1?

Doug Benjamin: I'm not sure — this was phrased as "prototype using PanDA to operate an inference service," and what you just described kind of already exists, maybe not at NERSC, but elsewhere.

Rui Wang: I don't know — let's make a note and come back to that.

Shawn McKee: Should I move you on?

Rui Wang: Please — yes, skip this one, go to the plans. Here are the plans — on the operations side, the top priority is finishing the NERSC allocation for FY26 and improving allocation efficiency. We also have DD allocations at the LCFs — low priority, just to keep them maintained for next year. On integration, we'd like to continue workflow testing — not just integration tests but scaling tests too, including integration testing of Globus Compute plus Harvester. We've actually renamed that to be the IRI API.

Unidentified speaker [Verena or Paolo]: I was about to ask — doesn't that seem like something worth putting effort toward? I think Jeremy might like that, especially if it's moving to 2.5.

Rui Wang: Yeah, we can add more detail — if you can send me the exact wording, or put it in a comment: should we list the IRI API specifically instead of "Globus Compute APIs"?

Unidentified speaker: Yes, put in the proper DOE terminology — IRI, or whatever it's officially called.

Rui Wang: I'd say so. Is that still priority 3?

Unidentified speaker: Priority 3 — meaning medium?

Rui Wang: Yeah, I think medium.

Unidentified speaker: Do you want it to be higher?

Rui Wang: I don't know that we can honestly say today it should be priority 1 — though I think it's vital for what we're doing.

Unidentified speaker [Verena]: I have a comment — I saw a similar one made later by Zach on another slide — that GPU and AI/ML workflow support shouldn't be marked low priority just because there's no formal request today. That doesn't mean the priority is actually low; it should be high priority, it's just not happening right now. Do you see what I mean?

Rui Wang: It sounds like — are you saying this should also be priority 3, on the next slide?

Unidentified speaker: I think so, yes.

Unidentified speaker [Verena or Paolo]: Native question here — how does this align with ATLAS's needs, especially regarding the Globus solution, Harvester, and Globus Compute?

Rui Wang: I think for Globus, it's approved by ADC. I'm not sure about IRI — that's just being approved, so we're planning integration tests, not only at NERSC.

Alexei Klimentov: My concern is — we don't have infinite personnel effort, and there's been a lot of discussion about this, at least in the last two HEP-CCE topics — Globus Compute API, Harvester, the Globus solution, and so on. This was discussed at CHEP, and experts told me it has nothing to do with what we're doing in ATLAS.

Rui Wang: What expert? Because for the most recent one, the data handling group specifically asked me to give a talk on this at the collaboration meeting — they wanted to know the status, and I didn't get any significant negative comments after the talk. Though someone may have said something out-of-band that I didn't hear.

Unidentified speaker [Verena or Paolo]: I think the question here is whether Globus Compute is the preferred way for leading facilities — local solutions probably aren't seen as a good option in their view.

Shawn McKee: I'm just thinking about the testing side of this — Ivan has his hand up, maybe he can comment.

Ivan: Yeah, I just wanted to comment — from the ADC/PanDA team's point of view, they confirmed they'd like to have the Globus plugin for Harvester implemented, and they'd welcome development from Rui on that. That's all I know, thanks.

Rui Wang: And Globus Compute — we already said should be replaced with the IRI API, right, here.

Doug Benjamin: Yeah, that's what we discussed — it may not be that different, but at NERSC specifically, we are definitely using Globus Compute.

Rui Wang: Should I go on to the next slide?

Shawn McKee: Please do. Alexei, are you happy with those answers?

Alexei Klimentov: Well, I'm not, but we can take it offline.

Rui Wang: Here are the new milestones from the plans — five new milestones for operations; only one is for NERSC, the others are for 2.3.3 integration.

Shawn McKee: The second one we were just discussing — making it medium — is everyone okay with that? It sounded like Alexei wanted the Globus milestone priority raised, but maybe I misunderstood — that's a discussion you'll have offline, and you'll adjust as needed.

Paolo: I have a question for you both — everybody else asks for more FTEs, but you two never do. Have you thought about what else you could do if the team were, say, 1.2 FTE instead of 0.7?

Rui Wang: Doug, do you want to answer that?

Doug Benjamin: I think this has to come from whatever's decided at the WBS 2.5 level, Paolo — it really depends on how big we want to scale up. Everything's money-limited, and we're trying to be good stewards of taxpayer dollars.

Paolo: Yeah, I don't object to that — I'm just saying, the HPC field is moving fast. I know we're getting a lot of help from outside ATLAS, through HEP-CCE — maybe that's the point to stress: that we're leveraging outside help, which is why we don't need a lot more labor. If we want to change the scope, that's a different story. For example, the workloads we're running — is it just full simulation, or fast simulation too?

Rui Wang: It's both full and fast simulation, yeah.

Paolo: CMS, if I'm not mistaken, also uses a Globus overlay cluster for everything — we don't?

Rui Wang: Right — we haven't gone that route. I don't know if that's something we should consider, I'm just saying.

Doug Benjamin: Yeah, think about it — there's an important thing we're not doing simply because we're only at 0.7 FTE. We'd do it, but we should also make sure Zach and project leadership support and buy into it.

Shawn McKee: So is that it — no summary slide yet?

Rui Wang: Fine, for now.

Shawn McKee: What's next? Okay, 2.3.4—

WBS 2.3.4 — Analysis Facilities (Wei Yang)

From this point, Rob Gardner's turns were explicitly labeled "Me" in the source transcript and are reproduced with only light cleanup. All other turns were labeled generically "Them," so the same inference approach used above applies to distinguishing Wei Yang (presenting) from Paolo, Verena, Alexei Klimentov, and Shawn McKee. One notable exception is flagged inline below: a stretch of dialogue about RAC meetings and Local Group Disk was labeled "Them" by the source transcript, but its content unmistakably reflects Rob's own experience and is attributed to him here — see Sources and Methods.

Wei Yang: Let me find where my slides are... let me share... oh, I need to go to that slide — here. I've got a slideshow, yeah, looks good — somewhere else, okay. All right, first one is the FTE count — I think there's a little change, but the SLAC number hasn't changed. The UC number — there's a request for some increase. BNL— sorry, we flipped over to slide two — are you on slide three right now?

Oh, okay — accomplishments, I should go there anyway, sorry. This is basically what we got from the monitoring dashboard we've been using for two years, so the numbers come from there, though there are some open questions. This is the number of SSH users, Jupyter users, batch users, and batch jobs/hours. For UC, we also broke out the percentage of users that are US versus international. There is a bit of a question about these numbers from BNL — if you look at this number, it's completely out of line compared to the other analysis facilities. I ran it through June 11, and yesterday I got different numbers for the same period. I think [Qiulan] needs to dig into the UC — sorry, the BNL — database to debug what's going on. I asked for a number from BNL — initially it came back as zero. She was supposed to fix it but hasn't gotten back to me yet; I asked her three days ago.

Unidentified speaker: Okay, don't pull everything in — okay, please.

Wei Yang: Yeah, anyway — I'm just pointing out these numbers probably don't reflect the actual number of jobs; I need the correct number. Otherwise, it's a relatively smooth operation — SLAC in particular is under significant resource constraints, but we've still managed to support our users. There are other accomplishments listed, you probably don't want me to read them all. On FTE: SLAC hasn't changed, UC is asking for an increase.

Unidentified speaker [Paolo]: Question about the table — what is the 63%, 78%?

Wei Yang: Yes — UC was able to break this down: 63%, for example, is the percentage of US users; the other 37% is international. Same for the other columns — percentage of US users versus worldwide.

Paolo: And the numbers in parentheses are last year's, I know — unfortunately they don't look very good for SLAC's brochure. It would be good if you backed up your statement that you're operating under significant resource constraints, because I could ask you right now: why on earth would I continue funding SLAC procurement if the number of users is going down? Well — wait, I didn't ask you that, but I could. What would the answer be? Presumably: we haven't gotten any new hardware, unlike UC, so users are moving there. But is there a way to show that BNL and SLAC's analysis facilities are saturated — that their smaller resources are basically maxed out?

Wei Yang: Yes — I don't operate Brookhaven so I can't answer for them, but for SLAC, complaints do come to me; one faculty member told me they moved all their users [elsewhere] because the resources aren't there.

Paolo: If you have any plots or tables showing that whatever small resources you do have are fully utilized, it'd be good to have in your back pocket — so when the question is asked, you can say "we're doing what we can." The same goes for BNL, I'd say.

Unidentified speaker [Verena]: Yeah, in general it'd be nice to see that kind of information — given the resources, how they're being utilized.

Paolo: Yes, that's actually a good point — even without the declining numbers, it would be good to see resource utilization.

Wei Yang: Part of the story is also that hardware gets retired and we're not buying new hardware to replace it.

Paolo: That's another good point — you guys should make these points, because otherwise it just looks bad. If you think that's the explanation, spell it out.

Wei Yang: Okay, let's move to FTEs. We already mentioned SLAC and UC. For Brookhaven, I believe the number is roughly correct, but there might be about a 0.1 FTE discrepancy in this table that I never called out, as I mentioned in my comments.

Shawn McKee: I think that's based on one of the milestones — but if BNL has made a change, we need to update this number and make sure it's consistent. Just a reminder for everybody: the goal was originally to get the effort spreadsheet correct first, and then have the slides follow the spreadsheet — that hasn't happened yet, and if you look at the current effort spreadsheet, these numbers don't match what's here. In fact, people need to be added — it looks like, Wei, you have Nakata in FY26, but there's no FY26 Nakata in the effort spreadsheet. Was he supposed to be added last year, or was he added? We need to make sure we're using the right numbers.

Wei Yang: I think the name was — that person was replaced.

Shawn McKee: Replaced, yeah, and I guess we never updated the name.

Wei Yang: You're right.

Shawn McKee: But Wei, you took a pass to make sure FY26 was consistent with Chuck's spreadsheet, right?

Wei Yang: Yes.

Shawn McKee: It still has the old name. My understanding was that was kind of gospel for FY26 unless there was a gross error — we should use whatever FY26 was; FY27 we can make changes. Generally, the spreadsheet should have been normalized first — everyone should have agreed on those numbers — and then the slides should reflect that. That hasn't happened in all cases, and I'm raising it as a concern, because when we get to scrubbing, people look at the various source spreadsheets, and if we don't agree, it causes a lot of discussion we don't have time for.

Paolo: Absolutely, 100% agreed — the spreadsheet is really important, because if it's not there, the money just isn't there. Remember, we're going to be under pressure — I'm not saying there will be pressure to cut, maybe there's pressure to increase, I have no idea — but we're going to be under pressure, and if we see a mess, we may look elsewhere. Be careful — it sounds like bureaucracy, but it's important bureaucracy. Same goes for the milestone numbers — we've asked repeatedly for the related milestone numbers to be added to the effort spreadsheet, just like you're doing here, but they're not there, at least they weren't there yesterday. We need that information, and you can't assume someone will add it for you, correctly, without checking. On the point of substance — we have, let's see, 0.35 FTE at SLAC, 0.5 or 0.4 depending on this Pelosi question, at Brookhaven, and 0.5 at Chicago, right? Is that correct? The user communities are very different, though — is there a scaling with the user community, or do we just need half an FTE everywhere, at every facility?

Wei Yang: There are two differences in what you're supporting — there's a base level of supporting any hardware and operating system that doesn't scale with users. But there are user-facing questions too, of two kinds: system-related ("I ran out of disk space," "I can't log in") and application-related, which I believe are mostly addressed by WBS 5 effort. It's not like everybody has the problem, but the more users you have, the more of this you'll see — though it won't double linearly. So there's a slow scaling factor, but it's not entirely obvious.

Paolo: Let me push on this a bit — do we have any cross-facility activities or milestones, where more than one analysis facility participates?

Wei Yang: You'll see, in terms of milestones, I'll list that later. In terms of support activities, that mostly happens at the physics-analysis level — for example, Shuwei will support some of the SLAC questions [WBS 5 effort], and he's not counted in 2.3.4. Within 2.3.4 itself, for things like

data transfers, we can give each other some help. But if it's related to the batch system itself, the other facilities probably don't know enough to help. Rob, I see [you have a comment].

Rob Gardner: I was going to say the one activity I can think of that goes across all three facilities is the benchmarking effort. I mean, this is not 2.3 specifically, but it's something that the 2.3 admins have to look after to support. But there is that thing that is, you know, I would say is a pretty standard thing that ties the three facilities together. I don't think that we have any effort available to do anything significant at the infrastructure layer, right, at SLAC and Chicago altogether. So those activities always sort of come as, you know, hope we can get to them, or hope we can do something. But it's hard to do that when you have so many daily events with users at your facility.

Wei Yang: As we saw, there isn't accounting sort of cross-facility, right? We tried to do the accounting — well, if it generated an idea, maybe this is the answer I had in mind — but we really, we call it analysis activities, but we are under 1.5 FTE [total], we have all this hardware, and the expectation is that we provide first-class support to ATLAS and the users.

Rob Gardner: Exactly.

Rob Gardner: Yeah, I agree. On a daily basis, you don't actually have a choice, right? When you have users pressing down on you for fixing problems and dealing with the infrastructure that's there — there's no choice in your to-do list for that day to do something more esoteric, cross-facility.

Paolo: So would this more esoteric, cross-facility activity be under the purview of 2.3.5?

Rob Gardner: Yes — and that's where we have had them, Paolo. And even there, the struggle has been getting effort from the labs — who can be assigned to work on those kinds of things. So we're limited there, mainly on the lab side, I'd say.

Paolo: On the lab side, at SLAC and BNL?

Rob Gardner: Yep. So, for example — you remember we had this manifesto during Lincoln's time to come up with a unified, one-stop-shopping place to log in, right? The challenge was finding the right layer that makes sense across the facilities. We backed off mainly because of the effort required across all three sites, and we've just had limited effort.

Paolo: Also, good luck convincing the labs to converge on a federated login scheme — IRIS-HEP is struggling to do that, imagine us.

Rob Gardner: Exactly. So we thought, okay, we have to choose the better opportunities — the things where you can make more vertical progress.

Alexei Klimentov: Okay, so the increase comes from who — which — yes, I remember. There's a change I don't understand: it looks like the total goes up by 0.2, which would be Chicago, but then there's an increase with Pelosi at Brookhaven that was never discussed at BNL. So —

sorry to repeat — that's not reflected correctly, because the total should then increase by 0.3, not 0.2. Earlier I commented it should be 1.35 — and that's probably what's in the spreadsheet.

Unidentified speaker [Verena]: Okay guys, sorry — just a quick thing, for Fengping: does this 0.5, in red in parentheses, mean that 0.2 moved to here and 0.05 moved to IRIS-HEP? Sorry, I was just trying to understand what it means.

Wei Yang: Oh, you mean this one — yeah, is this what happened, I think this is a previous number and then you moved — okay, this is the increase.

Verena: Okay, and one other curiosity — I see a lot of things listed here on different lines. Should some of these be in 2.3.5 instead? Especially on the second line — are those really things like ServiceX for operations, or is this more in the R&D category?

Rob Gardner: Yeah, it looks pretty lopsided there, doesn't it?

Verena: Yeah, sorry — I wasn't sure, or should we put some things in the other [category]?

Rob Gardner: If we wanted to rebalance a little bit, I do think we could remove some of these. I think ServiceX deployment on the AF is gated by [2.3.5.2], but we have additional support from IRIS-HEP for that, right.

Verena: It's also listed in 2.3.5.

Rob Gardner: So I think—

Verena: —2.3.5.2.

Rob Gardner: —the way we could probably subtract that out. I mean, the BinderHub, federated Kubernetes, notebooks, Dask, HTCondor stuff — I think you could take that off, because I've assigned that to Farnaz. Remember, Chicago's supported — not the infrastructure, the admin-type things Wei was talking about — it receives some bridging support from Midwest Tier 2 as well. So we have Midwest Tier 2 support coming in that we're leveraging — that's always been the case. That maybe partly explains why there's a larger portfolio of things there than at the others, because from the beginning we were only operating with 0.3 [dedicated AF effort], so obviously we're getting support from Midwest Tier 2.

Paolo: So — I wonder — this is a comment I also made about the 2.3.5 slides — throwing all this stuff here, what information does it give me? If I'm Paolo, or even Shawn, and I see all this here, what do I learn, to be able to judge? I think what's more important is to describe not the tools you're using, but what you want to use these tools for. It seems to be a general problem that we like to name cool stuff, and we're working on the cool stuff, but we forget to say why we're doing it — which is what I care about, in the end, to be honest.

Rob Gardner: I think that's a great suggestion, Paolo. I think the point of this slide isn't to go down any rabbit hole of some interesting technical thread — I think the point is to understand how the effort's deployed.

Paolo: And you can put in the comments — for the other sites too — note storage, notebooks, interactive usage, and so on, but don't go too much into detail. Okay, let's move on unless there are more questions.

Wei Yang: Okay, so for the challenges in FY26 — as you saw, all the first items for all the analysis facilities are hardware-related. But there are other things, like integration. And you see — another big item is GPU [demand] — this is a constant complaint that there's not enough GPU—

Rob Gardner: Comment here, under UC — next slide, before we get to this slide. I wouldn't open this door on this slide.

Wei Yang: Thank you so much. Okay.

Rob Gardner: I mean, I'd just delete that bullet — it's going to trigger a lot of unnecessary questions that wouldn't otherwise have come up.

Wei Yang: Okay.

Paolo: Also, I was going to — this sounds like a challenge, but I think it's an advantage. I was going to make a comment earlier on this: having four people at 0.1 each may not be the best way to manage a complex situation. So I don't see that as a challenge, actually — it may be a good thing. Anyway, the question of GPU demand — in particular, not batch GPU demand, but access to LLMs — that's the huge one, right?

Wei Yang: Yeah — so, GPU demand, I think the short answer is that at today's prices, forget about it. We really can't provide enough GPUs for users, so we have to rely on HPCs that have more capacity.

Paolo: I was just saying — I believe this is the complaint I received — is it because people can't migrate and do work at NERSC? There are already like 30 people actively working at NERSC now — I don't understand why people just — it's not convenient, right? But that's what people are doing.

Wei Yang: So I think that is — but for [LLM] models, it's another story. That requires — one thing is to provide allocation, like Ray was saying before — our support is to provide an allocation, that's easy to do. If we had to support running LLMs for the ATLAS community at ALCF, say, not always NERSC, that would require probably well over one FTE of effort, initially at least. So it'd be a significant enterprise. We may have to do that, but we're not ready to do that now, I don't think.

Paolo: Okay, challenges were well described. In terms of 2026 hardware: SLAC has asked for tape to back up data, and we're in a deployment phase. I don't know if Brookhaven has ever asked for or gotten any hardware. UC does have receipts, does have something deployed — this is completed, I believe. Rob, do you want to quickly go over it, or is it too much?

Rob Gardner: I don't have any comment. I don't know if you want to include this in your presentation — I don't know what others think, but this is just—

Paolo: The \$71K — where's that coming from? I'm curious about that, that's the only thing.

Rob Gardner: From our budget.

Paolo: From the AF budget, from the Midwest Tier 2 budget?

Rob Gardner: Midwest Tier 2 is probably how you'd label it, because we had a separate line item for the AF. It was a judgment call at the time — I think it was worth it.

Paolo: I'm not judging, I'm just asking — \$71K, so it's, okay.

Rob Gardner: I'd be interested to see what it would cost today — probably double.

Paolo: Yeah, that's probably familiar. Okay, fine, that's good. What else do we have? Anything else? Question — those [funds], are those the proposed funds for FY27, so not the current end-of-CA ones? These are the ones for the next CA — is that right?

Rob Gardner: No, I think this has already been spent.

Paolo: Oh, this is what you bought. Oh, sorry, okay.

Rob Gardner: Yeah.

Paolo: Okay, thank you — I'm sorry, I had a misunderstanding, thank you for clarifying.

Wei Yang: Okay, so — other ongoing M&O activities: normal things, maintenance, hardware upgrades, user support, monitoring, etc. In terms of milestone reviews, most of these are relatively technical and mostly on time or ongoing. There are two that we have to retire. The initial idea was built using WireGuard and Condor flocking — but when Lincoln left, we lost the key person to continue this. And this other one — we'd hoped we'd be able to do something about it, but the WBS 5 survey seems to indicate that users aren't interested.

Unidentified speaker: In those cases — which number are we talking about?

Wei Yang: This one, 537.

Unidentified speaker: Yeah, this is the one you retired because Lincoln left?

Wei Yang: Yes. And the other one, 538, is the one we'd hoped would let a user get remote access to the local group disk. But the survey by WBS 5 doesn't indicate that users are interested, so we're retiring that one too. But, as you'll see later, I still think we need to do something about it — there'll be a new, continued proposal next year.

Unidentified speaker [Verena]: So the first one — if you just said "I don't care," retiring it is the right answer. For the one with user interest, where you lost the person, rather than retire it you could change the baseline and say "we've now hired the new person, here's a new baseline" without retiring it. Is that what this is?

Wei Yang: Yes, that's right — this is the "bursting out" one. I think the idea also has some changes — in FY27 we're actually going to propose something similar, but different; this one will probably not depend on the same technology — we're probably thinking about using the IRI, or whatever it's called, the AmSC interface. But that does introduce broader uncertainties and risks.

Unidentified speaker: Okay — so you have not delayed the milestones, right?

Wei Yang: Well, these two — no, there isn't [a delay]. Maybe we move on. So, this is a summary of what we want to do in FY27. I'm sorry—

Rob Gardner: I was going to say, about that one milestone we retired because Lincoln left — remember, this was envisioned more than a year ago, and it really was kind of a first step toward some kind of federated platform between the three facilities.

Wei Yang: Both of them are actually bad ideas, [in hindsight].

Rob Gardner: And I think the winds have just blown differently this year, and our judgment has changed a bit on the usefulness of this, especially since there are other kinds of — let's call them facility API things — taking more priority, like AmSC. So—

Wei Yang: Yeah, that's exactly what I was going to say.

Rob Gardner: Yeah, so for me, I think the appropriate thing here is just to quietly retire it.

Wei Yang: Okay. Okay, so these are basically the things — an overview of what we want to give. We think about these priorities: the first is hardware — we hope there's money. The second, I do have a request, it's probably not realistic, we'll probably laugh, but I have a request: one big thing is, in particular, storage — we need to not just ask for storage hardware, but also put a storage policy in place, quotas and things like that; we'll have a milestone on this. Then we want to reorganize local group disk — again, this is 538 — because the pre-scrubbing slides show this might be underutilized, so we want to see if we can improve efficiency of using this space. And when users say they're not interested, I'm not sure it's because they're really not interested, or the PanDA/Grid commitment covers all their need, or simply because they don't know what capabilities we can provide. Then, utilizing HPC or AmSC resources for GPUs, especially batch — our previous idea was just simple Condor flocking to NERSC, but there are

other opportunities to explore. And last, we hope we can provide a general-purpose AI inference service, because we have grad students paying out of their own pocket to do these kinds of things.

Rob Gardner: Okay, so that's definitely a slide worth some thought. On point four, with the flocking thing — we heard from Rui and Doug that the GPU queues at NERSC are basically non-utilized. I know it's not quite the same thing as what you're talking about, but it's actually quite close.

Wei Yang: But yes — so couldn't that be the solution?

Rob Gardner: It could be — I will list it as a milestone, and note that I want to get help from 2.3.3, which is going to be [WBS] 2.5 [eventually]. Yeah, I guess my question is: should we distinguish access to LLMs versus other uses, like batch GPU training?

Wei Yang: Yes, they're separate. I'm thinking here about batch access to a GPU resource, to train a model — not about LLMs. LLMs are a different beast. On point five, which is connected to LLMs — yes, it's the simple solution, and I know it's not the perfect solution — it's simple: just pay, what is it, \$250 a year per student, and be done with it.

Unidentified speaker: \$250? It's—

Wei Yang: I think Claude Pro is \$20 — I don't know how much students would actually use, probably a lot, I don't know. I'm just raising this as a thing: do we want to go commercial? Do it ourselves? Use AmSC/HPC utility resources? I don't know the solution yet, but I want to bring this up and see if we have one. This is high risk, again, because we don't know what the solution is.

Rob Gardner: So, just to mention — there's NRP, for example, which provides access to open models on their platform, and Jordan has already figured out how to integrate it into a Jupyter notebook, so you can use the NRP LLM rather than a cloud subscription for reasoning against a notebook, that kind of thing. I don't know if there's an equivalent offering from the labs.

Wei Yang: There is — there's one from SLAC that they call I², I think, and there's one from Berkeley — how can I forget, I use it every day — I'm blanking on the name. There's an offering similar to that.

Rob Gardner: Is that the really nice one? I've been to the website, but it's only for Berkeley staff, right?

Wei Yang: Yeah, that's for sure — we'd need to negotiate with them if we wanted to extend it. But — thank you, Zach [for the reminder] —

Rob Gardner: —like, efforts in those directions, both of those, actually.

Wei Yang: Yeah, I suppose we need a harvester, whatever we can get, because you could have a situation where some users use a lot of this. By the way, this is one of those things which is at the edge between 2.3 and 2.5 — finding a solution may be part of 2.5. I think here we want to flag that there is a strong need for such a thing; where and how we solve it is a separate question. We should probably move on — interesting discussion though. Okay, so, these are the hardware requests — you'll love it, probably not realistic, but I put it here anyway. Brookhaven probably wants to modify these — I just exactly copied what I asked for.

Paolo: How did you come up with these numbers?

Wei Yang: Well, this is essentially a replacement — the disk replaces about 70% of what we have. I can ask for more if you have money.

Paolo: That's not the question. You know how complicated the modeling was for the Tier 2s — is the disk 10 years old, four years old, what is it?

Wei Yang: In terms of SLAC, everything is more than five years old — CPU and disk.

Paolo: CPU — would you have a complaint about 512 CPU cores?

Wei Yang: There's a lot — there's starting to be complaints. The disk is clearly close to six years old.

Paolo: My point is — not to discuss now, here — but you need to do a little more homework and come up with something justified. First, show your current hardware is saturated. Then, show the age profile of the hardware. Then, this is what we need to buy this year, in two years, in three years. For this kind of money, we need to do the exercise — we can't just come up with numbers like this. Okay.

Rob Gardner: Drop the URL in the chat showing just that distribution, Paolo, for UC — the specific number there is actually two-thirds of the cooperative-agreement year-one budget, that's how UC was defined.

Paolo: That probably did help — did UC do that as well, to get their numbers?

Wei Yang: So — FY27 plans and new milestones. A lot of these are more or less technical milestones; I'm not going to go through all of them. There's also some from SLAC and BNL — we probably need to talk through these again and make sure they reflect what BNL wants to do. These are the first 14 technical milestones. The last four reflect the overall plan across the analysis-facility level. One is that, even though the WBS 5 survey says local-group-disk access — part of our new idea of an integrated, federated analysis facility — isn't of interest to users, I'm not sure that's because users really aren't interested, or because the Grid middleware confuses them. So I want to shrink the scope: identify some space in local group disk — per the pre-scrubbing slides — and demonstrate we can quickly move a useful dataset in and out of that space, and document how a user can access it. There are risks: maybe users really aren't interested; and do we really have that much space in local group disk — the previous analysis

says yes, but we should verify again, implementing quotas and cleanup policies. I didn't put FTE numbers here because I think this will eventually be folded into the technical milestones at each AF — but I want to make sure the analysis facility overall does something about it, and that the policy WBS 5 sets up gets carried out. The last two are both GPU-related: can we utilize HPC or AmSC resources for batch GPU access — this needs us to work with 2.3.3, or in the future [WBS] 5. I put some effort here, though I don't know exactly how much we need — it depends on what resource and API we can find. On the AF side, I think it's relatively clear we want to provide a Condor-related interface to use the resource, but the type of resource is unknown.

Alexei Klimentov: Okay, so I'm a little confused. We just said you don't have the effort to do anything beyond creating accounts for users and dealing with emergencies. And now we have — 18, or fewer — what all look like R&D milestones to me. 90% of them look like R&D milestones, maybe it's 85%. So there's a clear contradiction: is this an operational WBS, or an R&D WBS?

Wei Yang: You can say this is for the future — what we need to do. Take this one as an example: we constantly get complaints from users about the analysis facility — what's the solution? If we don't do anything about it, all we do is keep hearing the complaints.

Alexei Klimentov: Well, but what is the ask? What are you asking — effort to do this? I don't see any effort listed here anywhere.

Wei Yang: Well, maybe in some of them there's 0.1 here, 0.2 there — with 1.2 FTE [total, across items].

Rob Gardner (*labeled "Them" in the source transcript, but the content is unmistakably Rob's — see Sources and Methods*): Let me also interrupt here for a second — Wei, I have attended the RAC meeting for the last 20 meetings, and I have not seen a single user complain about a shortage of space in local group disk in two and a half years, not one meeting. So I don't know where this milestone is even coming from.

Wei Yang: Yes — so this is also reflected in the pre-scrubbing slides, that local group disk is underutilized. So we have two choices for local group disk—

Unidentified speaker: I think it's full — you're saying it's full?

Wei Yang: It's not, it's underutilized — because we provision sufficiently for people so they can use it during their analysis.

Unidentified speaker: But now you're saying it's full, so we need to clean it — I'm just confused.

Wei Yang: It's a technical issue, but you should really bring it up in RAC — that's where we discuss local group disk, meeting every two weeks.

Rob Gardner (*again labeled "Them" in the source*): And I've been to at least 20 RAC meetings in the last two years — I've never seen anyone bring this up.

Unidentified speaker: You want to comment first? ... Rob's away anyway, okay.

Rob Gardner: I'm reading comments on milestones right now — sorry, new comments that have arrived since we started the meeting. Okay, what's the question? Sorry — are we talking about local group disks still, or—

Unidentified speaker [Alexei or Paolo]: All of this — most of it, other than [milestone] 16 — is more or less related to R&D: 15, 17, and 18. So is this really an operational milestone, or an R&D milestone? What do we want to get out of it?

Rob Gardner: I think anything that involves a new integration we haven't done before, off-the-shelf, you could classify as R&D. I think of the things we want to keep in 2.3.4 as just the operational stuff — no questions asked, no mystery. On the other hand, it's a place where you start to receive things being developed elsewhere — so the milestones we defined were really about bringing developments happening in 2.3.5 or 2.5 into a production space. There's a handoff between development and operation, and that's what was intended. That doesn't come for free, in other words — I don't think we have in mind effort in 2.5 operating AI infrastructure on our production facilities. If that's true, we can talk about it, but that's not what I understood it to be, and it's not consistent with the meeting we had with Kaushik, Shawn, and Zach — it was clear that the effort that would shift over to [WBS] 5 was more clearly in 2.3.5, which we identified.

Wei Yang: So — I guess I've raised the issue: we want to see if we can demonstrate local group disk, the GPUs — these are agreed as borderline R&D, and could happen somewhere else, we can remove them from here. But the issue doesn't disappear.

Rob Gardner: No, the issue doesn't disappear.

Unidentified speaker [Paolo]: The question really is — let's forget that, I'm just looking at the first slide. How many milestones are at 0.05 FTE? Are we looking at the right granularity of milestones? There's only one that's quarterly, which makes me think it's the kind of milestone I'd expect [as routine reporting]. So, again, I don't think this is serving you well, because it doesn't look feasible, even vaguely feasible — that's my problem.

Wei Yang: You mean the wording of my milestones? It doesn't look feasible?

Paolo: I mean, it just doesn't — forget even the last ones, the big problems you're highlighting as challenges. But putting a milestone with a September 2027 date on these big issues, when we don't know who's doing it or with what effort, doesn't work. It simply doesn't.

Unidentified speaker [Shawn McKee or Paolo]: Alexei, do you have a comment about this particular topic? I wanted to raise something.

Alexei Klimentov: I have a comment about this particular [issue] — we can discuss the particular milestone 2.3.4-15, but we have a fundamental question we don't know the answer to: what do our analysis facilities look like for High-Luminosity LHC? Everything else we're discussing is secondary to me, because for many years we thought wiser people in ATLAS or WLCG would solve this problem for us — it never happened. That's the fundamental question for me. After that, we can discuss whether a particular milestone needs 1 FTE, or 0.5, or several.

Paolo: Yes, I agree that should be the thrust, right? We need to understand what we think the AFs look like for High-Luminosity LHC, and then break down the particular milestones that support that vision — if it's not in the vision, we don't need the milestone. And [understand] what the community expects from us.

Shawn McKee: I had a different thing I wanted to raise. There hasn't been a good precedent for the milestones themselves — we've always had dates, but some milestones, all from 2.3.4, now say "ongoing" or "quarterly" in the date fields. I disagree with that. But if Paolo and Verena are happy to have those in there, it makes it difficult, because a milestone listed as "quarterly" or "ongoing" — I'm not sure when you report on it.

Paolo: Right, I guess you can do it — the way you do it for the Tier 2s works, right? You always assign dates.

Shawn McKee: And right now, in the date fields, there are recurring milestones which repeat every year, basically.

Paolo: That's right.

Shawn McKee: We have one year where we open an ongoing milestone, report on it at the end of the period, then make another one for the next period. So we need to fix that wherever it says "ongoing" and "quarterly." And Wei, when you made a copy, you made a copy of milestone numbers, and we now have duplicates in the sheet — 539, 540, 541, 542 are all duplicated. I can fix that, but the question is: should those have been new additions, or are they existing milestones?

Wei Yang: So — 539 through 542 — those are duplicates of ones that already existed in the spreadsheet?

Shawn McKee: Did you add them? Are they supposed to be last year's numbers, or are these new ones? The ones you have as BNL AF — are those new additions, or are they—

Wei Yang: No, no — these are from previous years. They weren't assigned a number before, and they weren't in the spreadsheet before — I checked Suzanne's original just now.

Shawn McKee: So you've added them somehow, but I'm not sure where they came from, and they're duplicating existing ones. If they are duplicates, I'll find them and replace the numbers — probably remove the old ones.

Wei Yang: Well, I'm not sure they were in Suzanne's original.

Shawn McKee: Anyway, this is a problem — so, just to get this cleared up: we should always have dates in the baseline and estimated fields, and watch out for duplicates. If they should have been in the original and are missing, we can do some spelunking to figure out what happened. Okay, sorry, that was a tangent, but I needed it.

Verena: That makes sense, thanks. So, I wanted to ask a couple of quick things. I'm a bit surprised — I was expecting to see a milestone for the UC AF for the installation of the end-of-CA hardware, potentially. Or is that not here?

Wei Yang: That's for next [year].

Verena: Is that somewhere? Did I miss it, sorry if I did? And the other question — I think, like in the Tier 2 presentation, we had a summary of the situation once we fold in the end-of-CA funds. I think that would be nice to have here too — because, for example, Rob, you shared this very nice, kind of frightening slide about the aging hardware and warranty expirations, and I think it would be interesting to see what happens once you fold in the end-of-CA purchases, since part of that is the refresh at the end of CA. I didn't see this information in these slides, if it's possible.

Rob Gardner: Just totally missed it, Verena — we should add both of those.

Verena: No worries — I'm sure Chris and John will be really interested in this.

Rob Gardner: Okay, we'll fix that.

Verena: Thank you. Also, final point — I'll repeat, I think there are two ways to tidy this up: you should consolidate, but when you consolidate, use the standard format, because tables like the one on slide 11 make it difficult to pick up the information — they should use the format Wei was using on slide 13, for example.

Wei Yang: Okay, let's move on — this is taking a little too long. So, risk: hardware is clear. Local group disk — whatever we do or don't do, there's a risk: we assume there's free space we can free up, but if there isn't, that's a problem. GPU is the biggest uncertainty — whether we do this or somebody else does, the resource itself isn't clear.

Unidentified speaker: So these are internal risks — you're not proposing updates to the risk spreadsheet?

Wei Yang: These are — I haven't updated the risk spreadsheet. I don't even know whether this is something we will do, actually.

Alexei Klimentov: Okay — again, I think there are a couple of risks presented here. Quick comment — when you're talking about GPUs for AmSC, I agree with the statement that for HPC you have it today. Why? Also, AmSC has several different APIs, as far as I can tell, even though the resource is the same — the API isn't necessarily the same.

Wei Yang: Move on, guys. Okay, I think that's about it — I'm not sure I want to go over all these milestones, but I do want to point out: you'll hear that the GPU/AmSC access is creating an imbalance among the AF users. I don't know how to address that.

Shawn McKee: Okay — I've tried to keep some notes; everyone else should feel free to add more in the shared notes doc.

Alexei Klimentov: All right, let me be brutally honest for a moment — we cannot go into the real scrubbing with these slides. I'm genuinely worried the reaction will be chaos at best, hostile at worst. We need to put together something proportionate to what we actually have. If we need to ask for more, it has to be justified — if you need 18 FTEs, ask for 18 with 18 milestones, not 1.2 FTE spread confusingly across three areas. Let's really think this through, not with our own heads, but with the head of someone unfamiliar with any of this, who's going to go crazy when we show them these things. So please, work on these slides. Also work on the effort spreadsheets and the milestones — they need work too. Please use this week, because right now these slides are, I think, dangerous.

Shawn McKee: Can I propose a five-minute bio break? Let's reconvene at 20 after — whatever that is your time.

Unidentified speaker: Yeah, 20 after.

(Break.)

Shawn McKee: Okay, so what's our plan to finish today? We have 2.3.5 and 2.3.1 left — both are big. There's no way we can do both.

Unidentified speaker [Paolo]: So can we move 2.3.1 to Wednesday?

Shawn McKee: Alexei, what do you think?

Alexei Klimentov: Either way — I'd rather we do 2.3.5 now, since I want to be able to give feedback on it.

Shawn McKee: So, Alexei, could you do 2.3.1 on Wednesday?

Alexei Klimentov: We haven't — yeah, and I need to create an Indico for it.

Shawn McKee: Other than Wednesday — are we overlapping with our normal facility coordination meeting, nominally at 1 pm? Should we do Tier-1 the hour before it, or—

Alexei Klimentov: Before probably won't work for me; after should be okay.

Shawn McKee: That would replace the facility coordination meeting that day.

Unidentified speaker: Sure, as long as nobody has anything pressing — we can just say 1 pm Wednesday for 2.3.1.

Shawn McKee: An hour should be enough, I hope — this has been a pretty useful discussion today, so we need it.

Shawn McKee: So, are we still on for 20 after, or do you need more time?

Ofer Rind: No, 20 after works — let's keep it short.

Shawn McKee: Thank you — go ahead, whoever's presenting.

Unidentified speaker: Assuming you can see that?

Ofer Rind: Yep.

WBS 2.3.5 — Continuous Integration and Operations (Ofer Rind)

Ofer Rind: Okay, so this starts with an overview of the activities across the four areas. A few things to point out: operational support is broken up into supporting, developing services, and facility R&D. And then there's this Level 4 — 2.3.5.4, cybersecurity — that was discussed at last year's scrubbing. It was implemented without dedicated effort allocated this year, but there was preparatory work done, which I'll detail.

Okay, this is the overall FTE summary table — I broke it up the way Shawn did as well. There's the effort shown, and the effort proposed to move into 2.5.

Listing some accomplishments — the main thing to point out is that Ivan [Glushkov] has been the ADC Coordinator for the past year, and we've done work on data challenges. There's also a lot done on ADC operations and developing services. Same for 2.3.5.2 — a big deal was completing the migration to Varnish for the conditions databases, and updating the ADC analytics platform; I'll get into that later, and I'll probably trim some of the detail here. A lot of work was done on agentic AI development by Ilija [Vukotic].

Unidentified speaker [Paolo]: Maybe you'll say more later, but in your previous table you had increased effort in agentic-ops development, but then also 0.6 going to 2.5. Can you explain the difference — why does 0.2 increase, and then 0.6 goes to 2.5? I don't understand.

Ofer Rind: The 0.2 is the increase — I think, Rob, correct me if I'm wrong — it's a shift of some of Ilija's effort, maybe from WBS 2.4, allocated to agentic development. That, plus the other effort he was already doing, totals 0.6 moving into WBS 2.5. So after this move, 2.3.5.2 will have 0.45 FTE remaining.

Rob Gardner: Yes, that's shown on the table I'll show for the Level 4 breakdown.

Unidentified speaker: Okay, okay, sorry — I got confused.

Ofer Rind: Yeah, understandable. A great deal of development effort went into facility R&D. Key things here — I'll summarize a bit more since it's probably too wordy — but getting Research

Platform RP1-core operational, and deploying the Open Data Facility at UC, were two very big efforts. A lot of the other effort here is in infrastructure/platform development and some agentic-ops work as well.

Unidentified speaker [Paolo or Verena]: Can you tell me what RP1 is?

Ofer Rind: That's Research Platform One. I forget, Rob — was this funded partly by IRIS-HEP, or is it still connected?

Rob Gardner: RP1 was the original Kubernetes platform that Lincoln developed. You'll remember it was originally designed to be a stretched Kubernetes cluster across multiple sites, and we pushed that about as far as it could go. A lot of this work was then picked up by Aidan Rosberg at Indiana this year — he took the platform development and made it more accessible for other sites to deploy. So rather than building one stretched cluster tied to a specific site, he developed what he calls RP1-core, which is basically an infrastructure-as-code pattern for deploying a Kubernetes cluster with all the features relevant to running an analysis facility.

Unidentified speaker: Is this specific to the analysis facility, or is it something other ATLAS sites have picked up?

Rob Gardner: It was actually designed to not be specific to our analysis facility — that's what the Open Data Facility is: a deployment of the RP1-core pattern as an independent platform. This converged with a lot of interest in open data — additional processing capability, additional mirroring of the open data we picked up this spring. So we thought: here's what we can do with RP1-core now that it's ready — we can take that pattern and deploy an ODF cluster with all the ingredients, so it could be done by other sites as well.

Unidentified speaker [Paolo or Verena]: I know you created a glossary, but what is Elwood?

Rob Gardner: There's a glossary at the end—

Unidentified speaker: You have more details on this later?

Ofer Rind: Yeah, maybe we'll return to it there.

Paolo: In general — as I said when I was commenting on the input slides for this deck — I don't think this is understandable by a normal management-level reader. It doesn't feel tuned for the audience. You should probably run it through an LLM and ask: how do I rephrase this for a manager?

Rob Gardner: I totally agree — after listening to the last presentation, and to all the presentations this morning, this is on me. I'll redo these slides as appropriate — make them readable for management, which means taking out all the jargon and so forth.

Paolo: Yeah.

Rob Gardner: So—

Paolo: Okay, let's do that together — I'll help with that.

Rob Gardner: Just as a heads-up — Paolo and I have been chatting, so there are going to be too many milestones as well; we already know we'll have to condense them.

Paolo: Yeah, okay — getting the message.

Ofer Rind: Okay, so this just shows that even though there's been no dedicated effort yet, preparatory work has gone on for 2.3.5.4, and some important things have been done. In particular, we now have a real work plan for addressing the gaps reported at last year's scrubbing. We've established the liaison model, made connections with OSG Security, and updated all the site security contacts. So this shows a great deal of preparation ahead of implementing the real plan. And I should point out — we actually started running real cybersecurity exercises.

Okay, now the tables — this is 2.3.5.1 Ops. A couple of things come up here. There's a small change — not really a change, but I've now shown, based on what Shawn indicated, some of the effort we proposed could move to 2.5, colored purple like he did. So out of the Ops effort, a quarter — this might still be under discussion with Armin — but Armen's effort is going to be allocated toward AI development work in 2.5. The other point to note — there was some question about Ivan's funding, whether it's US ATLAS or ATLAS. I think that's resolved: it's actually US ATLAS contributing this, wholly US ATLAS. I don't know if there's anything else we should add — Kaushik suggested something as well. And there's a small change for Shawn — his effort goes up on data challenges simply because DC27 is happening this year, which is a major effort before, during, and in writing the report afterward, and Shawn will be heavily involved. I think also — maybe this was an oversight last year — this was meant to be 0.2 and wasn't changed.

Shawn McKee: But I think it's easily justifiable—

Ofer Rind: Yeah, that was— yeah.

Shawn McKee: So we'd already agreed to move that last year, and indeed it was an error that it didn't propagate through to Chuck's spreadsheet.

Ofer Rind: And the other thing I wanted to say — this is basically reclassifying some effort that was in 2.3.6 that really belongs here, so it's not new total effort.

Unidentified speaker [Paolo]: So are you saying Shawn's effort is also reduced elsewhere?

Ofer Rind: Exactly — I think it's already been moved there, reduced there, but it wasn't updated here. That's all I'm saying.

Unidentified speaker: So is it showing in 2.3.2 — does Shawn have some effort there? It's not even showing him as present, working on that.

Ofer Rind: You mean this year, or for next year — or also this year?

Unidentified speaker: Backdated?

Ofer Rind: He backdated it to FY26. Yeah, it's complicated — but he was at 0.4 in Management last year and 0.1 here, and it really should have been 0.3 and 0.2 this year, FY26.

Unidentified speaker [Paolo]: Correct.

Shawn McKee: It's all right — my mistake, sorry. That move is from Management, not from Tier 2.

Unidentified speaker: Thank you — apologies.

Ofer Rind: So it ends up looking like 0.1 of new effort here — I'm not sure how to represent that.

Unidentified speaker [Paolo]: No — but please make a note, in parentheses, that it moved from — sorry, what's the Management WBS number?

Ofer Rind: 2.3.6.

Unidentified speaker: There we go — yeah, and the blue there is 2.3.6 at 0.3, and FY26 is 0.4. So that's where it moved from. You could just put: 2.3.6, FY26 0.4 → FY27 0.2, or whatever the right shorthand is.

Ofer Rind: All right. Any other questions on this slide?

(No response.)

Ofer Rind: Nope, okay. This one, we touched on before — in essence, Ilija goes from 0.8 to 0.4 in 2.3.5.2, with 0.6 as a candidate to move to 2.5 — from 0.8 to a total of 1.0, which is the sum of those two.

Unidentified speaker: Thank you — so this number shows the amount that would remain in 2.3.5.2. What was the rest of Ilija's effort before? Does anyone know? I think it was something in 2.4, maybe — Rob?

Rob Gardner: Yeah, it was 2.4.3, I think — that's where he was doing a lot of his initial AI-agent work, and also the chat-help development.

Ofer Rind: So it's really just moving it into different boxes here, not really changing the total. I think it'd help to make that clear — and we should double-check with the 2.4 folks.

Unidentified speaker [Paolo]: I can go take a look, noted.

Ofer Rind: Yeah — just put here that the 0.2 came from 2.4.3.

Unidentified speaker: Okay, that moved — all right, thanks, made a note of that.

Paolo: Because otherwise all these apparent increases raise a lot of questions — if it's just reshuffling, that causes far fewer.

Rob Gardner: In this case it is just reshuffling — and in fact, as you saw, I thought we had already moved all of Ilija's effort into 2.3.5.2 last year.

Ofer Rind: I thought so too, Rob — that's what surprised me to see.

Rob Gardner: I was kind of— but I stumbled over this a lot, right.

Unidentified speaker [Paolo]: That's also my understanding — and there's no Ilija line in 2.4.3 from last year.

Rob Gardner: So there you go.

Ofer Rind: Yeah, but he's not at 1.0 straight — I'm pretty sure this reflects the spreadsheet.

Unidentified speaker: No—

Ofer Rind: So I just checked Chuck's spreadsheet again, and it is split 0.8 and 0.2. Okay, so unfortunately it seems like it didn't get propagated through — sorry about that.

Unidentified speaker: Please continue.

Ofer Rind: Okay, Facility R&D — these are the people here, with a couple of changes. This isn't a change in effort, just a recoloring of some of it into 2.5 since it's AI-related. We lost Lincoln, as we all know, and [Giordon Stark] picks it up at 0.5 — he's 0.5 in WBS 5, with a quarter of that shifting to 2.5. So overall it's not much of a change: 1.25 down slightly to 1.2, with 0.5 moving to 2.5.

And this one — as I mentioned, we had zero dedicated effort last year for 2.3.5.4. We borrowed some effort from Shigeki, who does cybersecurity at BNL, and from Rob, to do the list of work I described earlier. But the plan has always been to have a funded FTE for this in FY27.

Unidentified speaker [Paolo]: Yep — Shigeki was supposed to spend, I want to say, 0.2 FTE — that's what we'd put in the planning for this. He's listed in 2.3.1 with effort on site security.

Ofer Rind: So I assume some of that time was spent on this, but it wasn't explicitly called out for 2.3.5.4. We never put any effort into 2.3.5.4 when we created it — it was always going to be funded by this new FTE. And all told, his actual contribution is a bit less than 0.1 FTE — that's

why I put a comment on my table not to forget it. I mentioned here that it's under 2.3.1.3, and it was less than 0.1 FTE, so I called that out explicitly.

Rob Gardner: Right, so I guess what I'm a little concerned about here is: should we be explicit that my effort is coming over here, or do I get to wash my hands of this?

Ofer Rind: So, yeah, that's about it — I'll add a bit more.

Rob Gardner: After describing — or am I going to be expected to contribute here further?

Ofer Rind: You mean in addition to this 1.0, or as part of it?

Rob Gardner: In addition to.

Ofer Rind: Yeah — there's the period after scrubbing, before we get the new money in the next CA. Do we just say we've got it ready and move forward with what we have, and no real effort happens in this area until we identify and bring on a hire? Or is there some interim effort needed just to keep things alive and make sure we're ready for the hire — some small setup work?

Rob Gardner: That, plus — the person who gets hired is going to need—

Unidentified speaker: Onboarding?

Rob Gardner: They're going to need onboarding and oversight — they're not going to be able to plug right into this and be productive on day one.

Ofer Rind: And also, the hire itself will take months. I believe there's an opportunity, and that's why we're asking for one FTE for the role. So a fraction from Rob and Shigeki will cover the gap, and then we'll see — somebody at, say, a quarter FTE.

Rob Gardner: I take—

Ofer Rind: But then, Rob — the idea would be that the planned FTE-in-effect for FY27 would be less than a full FTE, because of the time it takes to hire, and we're already partway into the year. So is there some effort we could ascribe to you and/or Shigeki to cover that piece? In essence, this 1.0 wouldn't be one person — it would end up split between you and whoever we eventually hire. The effort where we don't yet have the person to use it — we could shuffle some effort into that box in the interim. That's what I always thought we could do — what do people think?

Unidentified speaker [Paolo]: That's also an opportunity, depending on who you hire — you could split their effort across different 2.3 areas. This could be handled a number of ways; I don't think we need to decide immediately. What we know is that we put 1 FTE into the next CA, right?

Unidentified speaker: So Rob's question still stands, though — does he get to say "I've done my job for 2.3.5.4" and step away?

Ofer Rind: Did we put that in the next CA? Yes — I went and checked, it's a whole FTE.

Unidentified speaker: I thought we had only put in 0.2 — I thought we pushed for a whole FTE.

Ofer Rind: We put in a whole FTE, Rob.

Unidentified speaker: Oh, thank goodness — okay, that's good, I was very relieved.

Ofer Rind: Yes.

Rob Gardner: In the cybersecurity slides I created, I did talk about the kind of hire we'd want — somebody who isn't coming in as a compliance officer to hit everyone over the head with a security baseball bat—

Unidentified speaker: No, we don't want that.

Rob Gardner: Right — we want somebody coming out of our field, who can actually work and develop tooling, who understands what's coming in the agentic space in addition to traditional IT.

Unidentified speaker [Alexei]: Rob, I thought your slides nailed it for the kind of position we want. Let's move on, guys.

Ofer Rind: All right, challenges. I'm not going to go through this in detail, but support is a challenge — optimizing resource usage has been a bit of a struggle. We've been working with Taiwan, and there's actually been a great deal of improvement there, especially after the Taiwan site admin visited CERN for two months and worked with Ivan — so that's been both a challenge and a real accomplishment. We still have persistent issues with small sites and analysis users, that sort of thing. For 2.3.5.2, a lot of the challenges are around deploying the agents and the AI stack — I could probably consolidate this a bit. For 2.3.5.3, these are the changes in scope for agentic-ops development, actually starting to deploy these things now that users are starting to use them — there have been issues, some detail here on performance testing and security policies.

I put this here too — might move it to the extra slides — but since we had a hardware discussion: as most people know, UC operates the very large analytics platform for all of ATLAS. It's important to note this is a key data source for any of the agentic-ops work people are talking about, so it's important that it stays supported as that source. This describes the platform refresh, the hardware issues, and the details of what was spent — but this isn't asking for money, since it was already funded by the NSF CA end-of-award funds. I don't know if it's worth calling out here or just linking to it in the extra slides.

Unidentified speaker: I'm not sure this is a challenge, is it?

Ofer Rind: No, it's not a challenge — this comes after the challenges section; I guess it's the hardware section.

Paolo: One challenge in 2.3.5.3 — which I actually put in the comments on the wrong slide deck, sorry about that — is that I'd like to see a closer connection between what's described there, the milestones, and the new activities: essentially, what ATLAS priorities, coming from WBS 2.4, does this serve? That's an exercise the WBS 2.4 folks do all the time — they say, "we're doing this research because that's what the TDR is going to need." Honestly — probably because I can't understand what the "Shannon" thing is, or whatever jargon I'd need to navigate — I don't see that here. Even the slide I'm looking at right now feels like a collection of things we started working on for some reason. As much as we can, let's connect this to existing milestones and existing requirements, and anchor it — because otherwise it feels a little arbitrary. Okay, moving on.

Rob Gardner: Do you have a lot of comments about that, but okay, we'll move on.

Ofer Rind: All right, so that was ongoing activities — milestone review. Two are delayed: one because of delays deploying EL9, and a related item has had some delays too, but we expect it to complete soon. The other was delayed because Aidan still has plans to integrate with NET2 once people are available.

Now, plans — for 2.3.5.1 Operations, the main thing is that we're entering the Long Shutdown period, so we're looking at what operations will look like for HL-LHC and what needs to happen between now and then. We've been working with Kaushik and Ivan to identify operational issues where agentic ops might help, and we've been talking with Armin about that. The other big thing is DC27, coming up in the February–March timeframe, which will require a lot of concerted US effort to prepare for, participate in, and summarize afterward. Those are the main things for 2.3.5.1 in the upcoming year.

These are the proposed milestones — I put them in this table, but I'll convert them all to the format you asked for, Paolo. I may just list them and put the detail in extra slides like I did last year, but this summarizes it all on one page, hopefully usefully. It calls out the AI/ML-tools-for-ops effort that could move to 2.5, and the DC27 final report, which I wanted to highlight too.

Okay, this is where it gets a bit more new and jargon-heavy — 2.3.5.2 and 2.3.5.3. The idea is that a lot of work went into developing platforms and harnesses for agentic AI development this year, much of it by Ilija and others. To give you an idea, Paolo, it breaks into three pillars: the actual infrastructure and runtime harness for agentic work, which involves a governance framework — that's the "Shannon" piece, basically a way of controlling what AI agents do safely. That impacts both the operations pillar and the analysis-facility pillar — one improves operations, one improves analysis facilities, and both rely on this infrastructure. These efforts are, as noted, cross-cutting with 2.3.4, as we saw before. Hopefully that gives a clearer overview.

So in 2.3.5.2, this is essentially the new part — Track B is the agentic piece, which we're separating out to possibly move to 2.5. The traditional part Ilija's been working on is Andrew's [Hanushevsky's] services — caching, moving data, ServiceX/Y — and the analytics part, which

is very key. That's how it breaks up, and these are the milestones for it — a lot of these are the agentic work that could move to 2.5. For Track A, we don't have specific milestones right now, Rob — I think we need to get those for the core services. So that's the summary for now.

For 2.3.5.3, the effort you're asking about breaks into four thrusts, working across platform development and the back end. This is analysis-facility work, and this is the agentic work on the left — helping with monitoring operations and user analysis, where the LLM and the knowledge sets behind it come in. This side is mainly platform capabilities. It's broken out here what effort goes to 2.5 and how that would work.

Paolo: And these are— sorry, no, I— well, first, I'm still curious what that word ["Shannon"] is, but let me ask another question first. There seems to be — superficially, that's all we can judge from here — quite a bit of overlap between 2.3.5.3 and 2.3.5.2, both in personnel and in substance. Thrust one and Track B, for example, seem to be one and the same.

Rob Gardner: Right.

Paolo: And lastly — I don't know what it is, so I can't tell you if it's just relabeled.

Rob Gardner: So — Elwood is the user-space version of what Shannon is. You can think of Shannon as our path for introducing agents into our facility in a trusted manner, such that we have trust in what they're going to do. This comes from direct experience using AI coding agents on our AF — very easily you can break things if you have privilege to do so. So we realized we needed some kind of pass — some way of bringing agentic capabilities into infrastructure management itself. That's what Shannon is. Likewise, on the physicist-analysis side, it's a similar thing: we have users on the AF right now using AI assistants to do things, and they don't always use it wisely — sometimes they break things. Elwood is meant to bring facility awareness of what users can expect to be able to do on the AF, and give them a set of tools we know and trust that they can actually run as a user, in user space. That's the basic split between Shannon and Elwood. We actually have two harnesses running both versions of these agents: one privileged harness running on a dedicated box that runs NemoClaw, and another in user space with Elwood that handles user-level commands, where the security requirements are lower. This is the path we were driven to out of necessity, and we spent a lot of time developing the ideas more deeply during the Genesis Phase 1 proposal-writing process — I don't know that we have much chance of being funded by Genesis, but at any rate, we have the groundwork laid out for how we want to do this.

So, obviously, anything in the 2.3.5.2 space is where all the ADC-ops stuff goes — all the agentic work around CREST and caches, anything at the ADC-ops level that's of central interest to ATLAS. In 2.3.5.3, we're at the facility level, specifically focused on analysis-kind environments — that's the focus there. In that space, we don't have Ilija's effort in 2.3.5.3, but we still have to use, deploy, and test infrastructure — and that's the role I see RP1 taking. Both at the Shannon layer and the Elwood layer, RP1 is the first place it would touch metal, because it's an R&D platform: we can blow it away and rebuild it if something breaks, work through issues there, and then graduate them to their trusted counterparts on the production AF and

ODF. The idea is to use RP1 as the basic place where these things get developed, tested, and trusted, and then bring them over to our production environments where we actually have users.

Unidentified speaker: Yeah, it's just—

Rob Gardner: And very sensitive systems admins too, by the way — I didn't hear a lot about this at software and computing, but there's a whole other universe of introducing these things into very sensitive environments where there's decades worth of expertise in how to run facilities, and people aren't adopting at the same velocity — some people are very skeptical, some aren't engaging, others are more enthusiastic. That's another layer here that I think has to be looked at — that was one of the motivating factors for doing this: to basically create the system to gain trust.

Ofer Rind: Okay — let me try to summarize very simply what Rob said, Paolo: you're absolutely right that they look similar, because they use a similar framework and platform. The Shannon/Elwood system is something Rob and his team are developing that's applicable to other uses, other analysis facilities, and so on. In this case, one instance supports the analysis facility, on the RP1 side; in 2.3.5.2, it's used to support the regular Ops side — not the analysis-facility part, not RP1, but things like the "RobBot" triage agent he mentioned.

Paolo: More clear, or less? But yeah, I think this is very confusing. It was very helpful and super interesting, but the fact that it took 10–15 minutes of explaining — do you see what I'm saying? Maybe it's possible to simplify this a bit, or — I don't know, Rob, sorry, just—

Rob Gardner: Totally agree — and I'm very open to any suggestions on how to present this to a management audience. It's R&D, right, so it's going to have new things, because it's research. Any guidance you can give us to channel it, funnel it into a management review, would be very helpful.

Ofer Rind: We can give some suggestions. So — I have bad news, folks, because I just checked the NSF CA again, at the basis of estimate — we only put in 0.2 FTE for cybersecurity, and that's the 0.2 from Rob that we had planned. We hadn't put in any new additional effort. Unfortunately.

Unidentified speaker [Paolo]: That doesn't match at all with what I remember happening.

Ofer Rind: Well, let's go back and double-check the basis of estimate — but I think that would be good to clarify, because if that decision was made after we discussed it—

Unidentified speaker [Paolo]: I don't know — I was looking at the proposal, and I'm pretty sure it was 1.0.

Ofer Rind: Yes, but what worries me is that when the wheels hit the road, they put a different number into the actual budget. That would be very concerning.

Unidentified speaker: Okay, let's just double-check it, I'm sorry — I don't know, guys, you'll have to figure that out.

Ofer Rind: Yeah, let's flag that. Okay, moving on. So these are the milestones for that. On cybersecurity — the plan is to address these gaps by hiring the cybersecurity lead and establishing a functioning program, not just a plan on paper. I've put links here to exactly what Rob discussed earlier about what we're looking for in this person, and the details of the plans. There are about four slides of detail on people, which I kept out of here. Two things motivate this: first, the Trusted CI process exposed a number of gaps in our security posture, and of course we saw a site compromise last August. Second, in case people haven't been paying attention, the agentic-AI revolution is really changing cybersecurity at our sites — it's exposing a lot of different vulnerabilities, and all these new tools add new vulnerabilities too. So I think this is more important now than when we first started discussing it.

Unidentified speaker [Paolo]: Comment — I find boxes like the ones on this slide pretty dense and hard to take in while listening to a talk.

Ofer Rind: Okay, let's see if I can re-send it. So, these milestones reflect the work plan, and they reference the work-plan numbers in the extra slides. I haven't added these milestones to the spreadsheet yet, so I'm looking for guidance — I assume this will be a funded effort, but we don't have a person for it yet.

Unidentified speaker [Paolo]: The person to assign it to is just "cybersecurity lead" — should it go into the spreadsheet now, or does it need to wait for the hire?

Ofer Rind: No, it doesn't need to wait for the hire, but it needs to wait for prioritization.

Unidentified speaker [Paolo]: So what are the priorities here — are they all the same priority, and what is it?

Ofer Rind: Yeah, the priorities, I think, are listed in the spreadsheet.

Unidentified speaker: No—

Ofer Rind: These are in the spreadsheet, you put them in, Rob—

Unidentified speaker: Okay, all right, so they're in there — did you put priorities in the spreadsheet?

Ofer Rind: Okay, so we'll get the priorities in there.

Paolo: Yeah, my advice again would be to reduce the number of milestones unless they're really, really necessary. So — either these are very fine-grained, or you're not going to be able to do all of them. Do you have guidance, Paolo — sorry, do you have guidance on how much labor should be behind a milestone for it to exist? And how much labor is below the threshold where you just don't create the milestone? You understand my question?

Ofer Rind: I do.

Paolo: I used to say you wouldn't count 0.1 or 0.05 FTE at all in the activity spreadsheet, let alone as a milestone. A milestone should be a collection of several activities — I'd say at least 0.2 FTE. In a project with a couple-million-dollar budget, there's usually a handful of milestones per year.

Ofer Rind: But I was under the impression that everything we do has to have a milestone.

Paolo: It doesn't mean every activity has to be called a milestone — there could be one milestone that encompasses four activities.

Rob Gardner: I will take it upon myself to reduce this by a factor of 3.

Paolo: Thank you. As an example, you could have one milestone for each part of the work plan — so maybe four total, "implement this work plan," with details underneath.

Rob Gardner: We'll figure it out. But Paolo, please clarify about the funding source here — this is all predicated on having the FTE, right?

Paolo: We will clarify that, for sure. Even if there's no money, we have a plan — we may decide we want to do this rather than something else regardless. So, yes, let's clarify this well before — as soon as possible.

Unidentified speaker [Shawn or Alexei]: Maybe we can briefly discuss this on Wednesday — just double-check what the situation is, and if it really is as high-priority as we think, I might have to revisit other things and see what we can do.

Unidentified speaker [Paolo]: The other thing I wanted to ask — I thought there might be a part involving AI for cybersecurity, using AI to help with cybersecurity. Did I miss that?

Ofer Rind: Sorry, my brain is struggling — there's an entire work plan for agentic operations security.

Unidentified speaker [Paolo]: I think that's looking at the security of using agentic ops — I'm asking about the other way around: using AI for security.

Ofer Rind: I don't recall — Rob, maybe you remember? We were trying to do a Genesis proposal along those lines, but it didn't get funded.

Unidentified speaker [Alexei]: Maybe let's skip to the end, and I'll make a comment then — unless there's something vital you want to highlight in the risk analysis.

Ofer Rind: The only question I have — this is standard stuff, Rob gave me a bunch of risks here — I'm not sure whether this is the right format or level for these kinds of risks. I think we need to go over this, or Rob and I should look at it together to see if it fits.

Unidentified speaker [Paolo]: So — do you want risks here that are only the ones in the formal risk register, or also potential internal risks not in the register?

Paolo: Some of these are — I think it's okay to have internal ones, you can even put them in backup slides so you don't spend time on them in the talk, just noted as things you're aware of that don't rise to the level of the formal risk register. If there are things that should rise to that level, that's what we really need to focus on, because we'll need to add them.

Ofer Rind: Yeah — one thing missing here is a risk for 2.3.5.4. Based on what I'm hearing, that probably should have a risk: if we don't implement a cybersecurity lead, we don't have an incident-response protocol, we don't have evaluations of our security posture, and we could have sites taken offline due to cyberattacks.

Unidentified speaker [Shawn McKee]: Yeah, we should also look at the 2.3 Overview slide where I mentioned the risk I proposed adding — we could augment that one to encompass what you're describing.

Ofer Rind: Okay, all right — and this is just summary and conclusions; I call out the FTE here, mention DC27, and entering Long Shutdown 3. There's a whole bunch of extra slides you can look through, or ask about any of them.

Closing Discussion

Alexei Klimentov: Okay, so my general impression is that this was very useful, and I want to thank everyone for preparing literally hundreds of slides. One problem we have is that we have hundreds of slides and no time. Consider: if you're given, say, 30 minutes in the scrubbing schedule, you can typically talk for 15 minutes — maybe 20 if it's really boring and non-controversial — because questions come left and right, and I think we're asking fewer questions today than you'll get at the real scrubbing. Having listened to 2.3.4 and 2.3.5, I really don't see how you're going to fit into your allocated time budget. So maybe this is a discussion we need to have — either you reduce your presentation, or you're given more time, but I don't see how else this works. By the way, the time we had today doesn't represent the time we'll actually have in the scrubbing agenda — I've just put the link in the chat, the scrubbing Indico link, so you can see the time allotted for each area.

Unidentified speaker: Do you remember the password? Do I need to log in?

Unidentified speaker: If you log in, you can get in — my beautiful federated authentication.

Ofer Rind: I have 45 minutes, no problem — 45 minutes, 30 slides, probably. No, I'm going to shorten them — they'll get shortened.

Unidentified speaker [Paolo]: But the backup — more detail in the backup — the slides that did require further explanation, the ones actually shown—

Ofer Rind: Yes, those will be made clearer.

Alexei Klimentov: No, it's a little bit of an impossible task. Just to note — I did take notes for each presentation, they're in the shared doc; Level 3 leads should look at those for sure. It also points to Wednesday's pre-scrubbing for 2.3.1 — 45 minutes for 2.3.4 and 2.3.5 combined? We should fit, I think.

Shawn McKee: But I wanted to reevaluate, based on how the pre-scrubbing went, whether we need to adjust those times.

Alexei Klimentov: Well — I always try to stay positive, but I'm afraid that after Wednesday we should schedule one more — not a rehearsal exactly, but a session between 2.3 people, just Level 3 leads, you and I, to go through the slides, because I have three pages of notes. I think the timing is okay. I was very, very worried about the 2.3.4 slides, I already said that. The 2.3.5 slides are better, but yes, we need to work on these slides. My impression — I want to be very direct — is that for 2.3.4 and 2.3.5, we did not deliver the message. We described minor activities but didn't deliver the message that matters. I expect we'll get a bunch of comments on Wednesday. That's why I'd suggest we find time on Friday for 2.3 to go through the slides together.

Shawn McKee: Okay, sounds totally appropriate to me — people should look at their schedules for Friday and try to keep it as clear as possible.

Rob Gardner: I mean, just to push back a little bit, Alexei — for 2.3.4, we have one-point-something, 1.2 FTE split over—

Alexei Klimentov: Exactly, exactly, Rob — but when exactly—

Rob Gardner: —supporting more than 500 active ATLAS users.

Alexei Klimentov: Exactly — but this question, which was asked today, has already come up twice: what is our ultimate goal? Is it just to have a fraction of people supporting users and not pretending to do anything else? Or do we integrate with 2.3.4, 2.3.5, etc., and ask for at least one more FTE, say — or collaborate with 2.5, if the key effort comes from 2.5? That's the point. Yeah, I agree with you — with 1.2 or 1.1 FTE, you can't do anything else.

Rob Gardner: But my comment wasn't meant to be negative — I was asking, hoping to turn it around: based on what you've seen and what you know, what's the best message that's going to carry weight for the management team? Part of this is maybe talking to Viviana about something — I'm not sure the survey quite captured it, but I think there's a certain kind of user experience that's happened over the last year that we could capture. We don't do things like collect success stories — like the large HPC centers do, where they pester people to tell them what great new science their allocation achieved. In ATLAS, we don't do that. But if we had something—

Alexei Klimentov: But a success story can come from those utilization plots Paolo asked about earlier — for each analysis facility, demonstrate uptime and utilization; if you have 100% uptime and 100% utilization, what more can you do? But we need this discussion — I saw you touch on

it just now. We also have to answer, again, why three [Analysis Facilities] — reviewers will keep asking that.

Rob Gardner: He always asks that question.

Alexei Klimentov: And this question — why we need to distribute analysis facilities across three sites — but for R&D, I'd say we need a plan: we can say, okay, all our engagements will be reviewed in a year, and we'll stop half of them and continue half. That's fine — not every one continues forever, some will be reviewed and continue, some won't. This is a discussion I'd suggest you have for 2.3.

Rob Gardner: Just — thinking about the three facilities, in a way it's mitigating risk. You're providing an analysis resource at three different places with different constraints, different hardware, different potential issues for creating problems or losing capacity. It's about mitigating risk and maximizing opportunity — the fact that you have an Analysis Facility at Midwest Tier 2, co-located with Midwest Tier 2, means you can flock jobs into idle Midwest Tier 2 cores from the Analysis Facility, for people doing that kind of thing. These are just suggestions.

Alexei Klimentov: In this case, it contradicts what you showed — again, I want to be very direct: if you see a downslot, and demand can't cover it—

Rob Gardner: Yeah, I was thinking of it more programmatically, not on a daily operational, time-sensitive level. If a site were to go away for some reason, you have an experience base of analysis support and analysis facilities at two other sites to try to pick it up.

Alexei Klimentov: How do we share our expertise across those sites? We don't. And can we use the same technology everywhere?

Rob Gardner: Well — we kind of are, in some cases, not for every service, because there are constraints.

Alexei Klimentov: But, for example, if we decide the best storage-organization solution is some specific technology — pick your favorite — should we deploy it everywhere?

Rob Gardner: RP1 is, of course, an example of trying to make something that should, in principle, be extensible to other sites — but we'd have to demonstrate that it gives a functional basis to do everything an analysis facility needs to do. So maybe that's a direction to pursue, I don't know. But I think we need to be able to answer the question when it's asked, because we know it's probably going to be asked.

Shawn McKee: Okay, all right — so we have the notes, lots of stuff there. For others who've taken notes — we'll certainly discuss, but do you want to add them to the common notes doc, mail them out, or selectively contact people with your comments?

Rob Gardner: Shawn, I've been taking notes in Granola, so I can post those into the common notes there.

Shawn McKee: Okay, perfect — all right, so I guess we'll leave it there, folks. Alexei, can you stay on for a minute? We wanted to chat.

Alexei Klimentov: Okay.

(Closing chatter, not fully intelligible.)

Shawn McKee: Bye, everybody, thank you.

[Transcript ends here.]

Sources and Methods

Inputs provided:

1. A cumulative Granola transcript of the meeting, supplied in three installments (each later one not trimmed from the earlier ones). For the Overview, 2.3.2, and 2.3.3 portions, no speaker diarization/names were included — Granola labeled all turns generically as "Them." Starting with the WBS 2.3.4 portion, Rob Gardner's own turns began being labeled "Me" in the source transcript (rendered here as "Rob Gardner," per Rob's request to make the record general for the whole team rather than framed around a single participant's perspective), which meaningfully improved attribution from that point on. The 2.3.5 and closing portions continued with the same "Me"/"Them" pattern.
2. Rob's verbal speaker hints, provided across the three installments: for the first three talks — the Overview was given by Shawn McKee; 2.3.2 by Fred Luehring; 2.3.3 by Rui Wang and Doug Benjamin; "Paolo" and "Verena" asked most review questions; Alexei Klimentov asked about Globus/Rucio and separately objected to the scope of Globus-related work; Ivan made one identified comment; Paolo asked "why don't you ask for more FTEs"; Doug Benjamin responded. For 2.3.4, Rob indicated Wei Yang (SLAC) presented, and flagged that the presentation had a number of problems, some requiring Rob's own follow-up. For 2.3.5, Rob indicated Ofer Rind presented, asked that action items be split per WBS area with a general list at the end, asked that "Me" be replaced with "Rob" throughout, and confirmed that Paolo was the one consistently urging the group to keep moving through the slides.
3. Live content from the five relevant Google Slides decks in the "2026 WBS 2.3 Scrubbing" Drive folder (Overview, 2.3.2, 2.3.3, 2.3.4, 2.3.5), pulled directly to cross-check figures, milestone numbers, and terminology referenced in the discussion. The 2.3.4 slide data confirmed the FTE-total arithmetic issue (1.05→1.25 excludes Pelosi's new 0.1 FTE) with reasonable confidence. The 2.3.5 slide data (FTE tables through the Challenges slide) confirmed the FY26/FY27 numbers for Ofer Rind, Armen Vartapetian, Ivan Glushkov, Shawn McKee, Ilija Vukotic, Aidan Rosberg, Giordon Stark, and the cybersecurity new-hire line, and made it possible to correct two names Granola garbled: "Armen" was transcribed as "Alec," and "Giordon Stark" was transcribed as "Jordan." The 2.3.5 deck's milestone-review, risk, and summary slides were not fetched

(only the first 12 slides loaded in the session); the transcript's own detail was used for those sections instead.

Labeling anomaly in the 2.3.4 section: two passages — Rob's comments about having attended roughly 20 RAC meetings without hearing a Local Group Disk complaint — were tagged "Them" (the generic non-Rob label) in the source transcript rather than "Me," even though the content is unambiguously Rob speaking in the first person about his own committee attendance. This is most likely a Granola voice-diarization error (e.g., a moment where audio routing or cross-talk confused the speaker model) rather than a real second speaker. Both instances are reproduced above attributed to "Rob Gardner," with an inline note flagging the discrepancy, rather than filed under "Unidentified speaker" or left mistakenly attributed to Wei/a reviewer.

Method: The transcript was corrected for phonetic and garbled errors, broken sentences were completed where meaning was recoverable, and consecutive utterances from the same identified speaker were combined into unified paragraphs. Because the source transcript carried no speaker labels at all prior to Rob's own turns being marked, attribution beyond the explicitly hinted exchanges is inferred from context — direct address ("Shawn, feel free to...", "Thank you, Paolo"), continuity of a speaker's own slide narration, and consistency with the technical content each named individual would plausibly own. These inferred turns are marked "Unidentified speaker" with a bracketed best guess where the context points fairly clearly to Paolo or Verena; genuine uncertainty is preserved rather than guessed away. A few short, non-substantive fragments (false starts, filler, a few indecipherable phrases) were dropped for readability without altering meaning.

Significant inferences and corrections worth flagging to Rob:

- Per Rob's correction, Paolo was consistently the one urging presenters to keep moving through the slides rather than reading them verbatim. Two early instances of this pacing role (telling Shawn to "fly through the slides" in the Overview, and the "we're not reading the slides verbatim" reminder in the Tier-2 section) were initially misattributed to Alexei Klimentov and Shawn McKee respectively in an earlier draft and have been corrected to Paolo throughout.
- "Serene," heard in Fred's aside about "wrangling with [X] at length" over the funding slide, is a phonetic best guess for a name; it was rendered generically as "the reviewers" in the cleaned transcript rather than guessing a specific person. (Srini)
- Full names/affiliations for "Paolo" and "Verena" were not stated in the recording; they appear in the metadata and transcript exactly as given.
- In the 2.3.5 section, "Jordan" (as transcribed) is corrected to Giordon Stark (UCSC) based on the slide deck's Facility R&D FTE table, and "Alec" is corrected to Armen Vartapetian based on the 2.3.5.1 Ops FTE table.
- A handful of turns across the 2.3.4 and 2.3.5 sections are attributed to "Shawn McKee" based on strong stylistic/topical continuity with his effort-spreadsheet-as-truth and milestone-bookkeeping role from the Overview, rather than an explicit name anchor in

those specific exchanges; flagged as a slightly lower-confidence attribution than the anchored ones.

- The transcript now runs through the full session, ending as the meeting closes; WBS 2.3.1 (Tier-1) was not covered in this session and was rescheduled (see Meeting Metadata).

WBS 2.3.1 Tier-1 (Rescheduled to July 8, 1 PM Eastern replacing Facility Coordination meeting)

We will use Indico / Zoom for Facility Coordination: <https://indico.cern.ch/event/1673148/>

Alexei presented, skipped to slide 5 at Paolo's request.

Paolo: Question about disk delivered vs what is total disk available at Tier-1?

Fix total on "blue box" storage amount in comment below: 67 PB -> ~77PB

\$27K discrepancy on VHW numbers (\$650K on detail slide vs \$677K on summary)

Notes from Rob (Granola + Claude) Wednesday

2026-07-08 US ATLAS Facility Coordination — WBS 2.3.1 Pre-Scrubbing Review

Meeting Metadata

- **Date:** Wednesday, July 8, 2026
- **Title:** US ATLAS Facility Coordination — WBS 2.3.1 (Tier-1) scrubbing slides pre review
- **Main presenter:** Alexei Klimentov (BNL, WBS 2.3.1)
- **Participants (inferred):** Alexei Klimentov, Paolo Calafiura, Verena Martinez Outschoorn, Shawn McKee, Ofer Rind, Rob Gardner, Fred Luehring; likely Qiulan Huang and possibly others
- **Context:** Prep for the annual US ATLAS scrubbing (Alexei presents July 14, 2026, at Stony Brook)

Reference Links

- Attached slides: *WBS 2.3.1 2026 Scrubbing Slides* (A. Klimentov, July 14, 2026) • FY27 effort spreadsheet (referenced throughout; URL not stated in transcript) • US ATLAS risk register: <https://www.usatlas.org/risk-management> (cited on slide 43)

Summary

This was a pre-scrubbing walkthrough of Alexei Klimentov's WBS 2.3.1 (Tier-1) slides ahead of the July 14 scrubbing at Stony Brook, with Paolo and Verena asking most of the questions and additional input from Shawn, Ofer, Fred, and Rob.

Much of the first half concerned personnel and the apparent mismatch between the slides (10.75 FTE FY26, 11.25 FTE FY27) and the draft effort spreadsheet (10.59 / 11.0). Alexei explained the quarter-FTE discrepancy through hiring timing rather than scope changes: N. Answer was hired to cover work formerly done by Latif (who left BNL in July 2025) at substantially lower cost, with the savings also funding part of the replacement for Garcia, who retired (search reopened after the leading candidate's credentials failed HR verification

— 0.05 FTE FY26, 0.15 FTE FY27 planned, within budget). In storage, the “Spradley replacement” was covered part-time by A. Undrus and Z. Yang before S. Huysman was hired in June 2026 (0.2 FTE FY26, 0.5 FTE FY27). Alexei is also proposing a new 0.25 FTE storage R&D effort for FY27 — entered in the effort spreadsheet as TBD — which Paolo and Verena endorsed but asked to be justified with a few bullets, since it falls under 2.3.5 R&D and all changes will be scrutinized at scrubbing. It was clarified that Misawa's 0.5 FTE includes the cybersecurity fraction now organized under the dedicated 2.3.5.4 area. Overall personnel budget remains below FY25, still not recovered from the FY25 cut.

On hardware, Alexei walked through the market-trends slide (memory up ~5x, SSDs up ~5x, HDDs up 30–50%) and the storage age-profile table: the FY27 disk pledge is 49.5 PB, the current disk forecast is 67 PB, and if the 16 PB of 2021-vintage disk is retired the pledge cannot be delivered. The 10.6 PB “early purchase” was made, but the 6.4 PB second tranche (~\$640k at the May estimate, possibly ~\$700k now) of what was originally a single planned

1

FY26 purchase was postponed; the group agreed to verify with Srini whether that money is still available, since ordering now would still mean December delivery. Discussion confirmed the bulk-storage policy remains HDD (NVMe/SSD only for specialized needs), which argues for prioritizing storage purchases while HDD inflation is comparatively modest. Fred flagged an internal inconsistency (67 PB forecast vs. 77 PB with the 10.6 PB included) that Alexei will fix, and Paolo asked for a pledge column next to the age-profile totals. The data-center power instabilities were reviewed: PDUs manufactured during the pandemic and aluminum cabling are design-level issues that cannot be fixed, while maintenance MoU gaps and generator switchover procedures have been addressed; a planned power upgrade (to 4.8 MW) comes in July–August 2026, paid by the lab with US ATLAS paying only its recalculated fair-share fraction of data-center power (that recalculation, plus bulk installation, is why the power line item dropped ~\$100k).

The budget discussion took time to untangle but landed clearly: against Chuck's FY27 guidance of \$3,965k (of which \$1,737k capital), the FY27 capital forecast is \$2,128k. FY26 shows a \$442k operations surplus (real economization, achieved transparently without moving money

between operations and capital) against a \$793k capital deficit — deferred purchases include the \$450k virtualization-hardware refresh (only a \$75k warranty extension was bought) and the delayed disk. With the FY27 shortfall of \$391k, the accumulated hole is over \$1M, on top of a \$1,048k cut in August 2025 and ~\$1.2M in the prior two years. Rob commented on the virtualization hardware costs (\$650–677k for a 100% refresh in FY27), noting that given this year’s price changes and that these are likely premium NVMe drives backing the VM system, the figure is probably not unreasonable. Paolo, initially “stunned” by the numbers, concluded the presentation must make explicit that the Tier-1 has survived successive savage cuts, has accumulated a debt of deferred procurement, and has essentially nothing left to cut; he will propose a concrete reorganization of the surplus/deficit presentation.

On milestones, four FY26 milestones are delayed, all tied to the personnel and hardware cuts in storage. Coordination with the new WBS 2.5 (AI/ML) was deferred to in-person discussion at Stony Brook next week — Kaushik and Zach need to meet with Shawn and Alexei to identify joint milestones and possible FTE swaps; nothing needs to be finalized before scrubbing since milestones will be updated afterward. Shawn closed with logistics: a Friday 1:00–3:30 pm ET session to run through all L3 slides (Alexei will present his final set); a request that Wei, Ofer, and others look at the less-detailed 2.3.4 slide format Rob circulated that morning as a possible convergence point; a newly identified problem in 2.3.5.4 cybersecurity — the next cooperative agreement contains no new money for that effort, and the 0.2 FTE appears to have been taken from elsewhere in 2.3.5, which needs discussion; and a plea for everyone to re-verify the effort spreadsheet without moving anything to 2.5 yet, since the budget rollup for Srinu is already behind. It was agreed that FY27 proposals should show full effort side-by-side with the proposed 2.5 moves, with the actual move applied afterward, so nothing falls through the cracks. Ofer raised whether the power-instability slide should show the job-monitoring plot (throughput down ~30% during the incidents even though pledges were delivered), and flagged that Louis’s containerization work — done mainly for the traineeship Verena runs, not as a 2.3.4 deliverable — was mis-mapped to a 2.3.4 milestone; a dedicated 2.3.4 discussion is set for Friday.

2

Action Items

Owner Action Context / Deadline

Alexei Add justification bullets for the proposed 0.25 FTE
FY27 storage R&D effort;
share with Ofer (2.3.5) and
keep in back pocket for
scrubbing questions
Alexei Clarify in slides why FTEs did not
increase (late
hires/replacements: Anwer,
Garcia replacement,
Huysman) vs. effort
spreadsheet
Alexei Fix 67 vs. 77 PB inconsistency in
disk

forecast; add pledge
column (or verbal note) to
storage age-profile table
Alexei / Shawn Verify with Srinu whether
the \$640k (~\$700k current
estimate) for the deferred
6.4 PB disk purchase is
still available
Paolo Draft concrete proposal to
reorganize the budget
presentation
(surplus/deficit framing;
accumulated procurement
debt; “nothing left to cut”)

Shawn Revert his effort-spreadsheet edits that pre-applied the 2.5 moves (comments in the 2.3.5 FY27 rows)	3	Before July 14 scrubbing Before
All L3 leads Double/triple-check the effort spreadsheet; do NOT move effort to 2.5 yet — budget rollup for Srinu is behind		scrubbing
Rui (probable), all affected Show full FY27 effort side-by-side with proposed 2.5 moves in slides; apply moves after scrubbing		Before scrubbing
		Urgent — December delivery even if ordered now
		Before scrubbing
		Now
		ASAP
		Before scrubbing
Owner Action Context / Deadline		
Kaushik, Zach, Shawn, Alexei	In-person discussion at possible FTE swaps Stony Brook on joint 2.3/2.5 milestones and	Next week
	less-detailed 2.3.4 slide format circulated July 8; converge on level of detail	
Wei, Ofer, others Review Rob's	Ofer Follow up on Louis's effort attribution (containerization work)	

mis-mapped to a 2.3.4 milestone; 0.2/0.1/0.1 in FY26)

Ofer / Alexei Decide whether to include job-monitoring plot showing ~30% throughput reduction during data-center power incidents (slide 27)

Shawn / Rob Address 2.3.5.4 cybersecurity funding gap: no new money in next CA; 0.2 FTE taken from elsewhere in 2.3.5

Verena Respond to Shawn's email re: Ilija's effort (2.4.3.2 question) Before Friday

With Alexei/Qiulan Before scrubbing

Open — needs discussion Sent July 8

Alexei, Shawn (+ optional others)

WBS 2.3 L3 slides, incl. Alexei's final set and dedicated 2.3.4 discussion Friday, 1:00–3:30 pm ET

Transcript (cleaned)
Friday run-through of all

Speaker labels are inferred from context — the Granola source labels only “Me” (Rob) and “Them” (everyone else). Uncertain attributions are marked (?).

Fred Luehring (?): So there's Answer — that's Tier-1 operations and monitoring — and there's also the replacement for Garcia, since Garcia retired. That's slide 22, I believe. On this particular slide, Alexei, did you make it clear that these don't match the effort spreadsheet because of the new-hire situation?

Alexei Klimentov: Infrastructure matches the effort spreadsheet. Fabrics matches. Storage does not, for two reasons. First, we had what was called the Spradley replacement: A. Undrus

4

and Z. Yang worked in the storage team part-time before we hired Stefan Huysman, so you don't see their effort in the spreadsheet, but it was used for storage. Second, there is a proposal for a quarter-FTE for storage R&D. I don't claim it must sit in Tier-1, but I claim it is important.

Shawn McKee (?): The reason I made it yellow is that you don't have a name there.

Alexei Klimentov: Correct, we don't have a name yet.

Paolo Calafiura: One thing we should agree on, because it may cause confusion: in these slides you are not showing the ideal, long-term situation. If Garcia left and somebody replaces him, I would expect the FTEs not to change. But you changed some because people were not available or not yet hired. That could cause confusion.

Alexei Klimentov: Slide 22. I should also mention — I always forget — that a fraction of Misawa's effort is for cybersecurity, but we still count it here.

Paolo Calafiura: And cybersecurity — is it part of infrastructure or fabrics?

Alexei Klimentov: We agreed it will be a dedicated level under 2.3.5 — Rob and Shigeki worked on this — but it is still counted here for now.

Paolo Calafiura: So Misawa's 0.5 includes cybersecurity. Let me make a note of that so we don't get confused later.

Verena Martinez Outschoorn (?): And Alexei, just to understand — the storage R&D is not Misawa; it would be another person?

Alexei Klimentov: Another person. And if we agree we need storage R&D, in my view it should be a bit more than 0.25, done together with the universities, and probably it should go under the R&D area — but I wanted to bring it here, because we did work on this.

Verena Martinez Outschoorn: I think it would be great to have a couple of bullets for Paolo and me to make the case. Storage R&D, with what we've been seeing recently, is ever more important — but all the changes are going to get grilled, Paolo and I for sure.

Alexei Klimentov: I'm happy to do it. I can prepare it and give it to Ofer, because R&D is 2.3.5.

Verena Martinez Outschoorn: You have it here, so just one or two bullets to keep in our back pocket. Two names came up several times: Garcia and Anwer.

Alexei Klimentov: Garcia retired and we are in the process of hiring a new person. Anwer replaced Latif, who left BNL in July 2025. Latif was a very expensive person; Anwer is doing the same job with a bit more FTE but at a fraction of Latif's cost — that was one of the agreements with Chuck when we hired him. And to be clear, the person doing the operations-manager job now is John McCarthy — McCarthy was always in the list; we did not increase his fraction. So it is really two new hires: Anwer, and the Garcia replacement — and the total stays below the budget we used with Latif.

Paolo Calafiura: And Garcia — do you have a search open?

5

Alexei Klimentov: The search is open. We had a leading candidate who received an offer from BNL, but HR then found that some of his diplomas were not correct, so the search is ongoing. We are interviewing.

Paolo Calafiura: In the long run you expect 0.15 of the new hire, and to stay within budget. That's clear, thank you. And at a very high level — is the personnel budget changed with respect to last year?

Alexei Klimentov: We have Chuck's projected annual increase. But compared with FY25, our budget will be less than FY25. We have not fully recovered — we had a severe cut in FY25.

Verena Martinez Outschoorn: My only point: the draft spreadsheet right now says FY26 is 10.59 and FY27 is 11.0, and this slide shows 10.75 and 11.25 — so there's a quarter FTE to understand. Going through it: the 0.1 increase for Anwer, the 0.1–0.15 for the Garcia replacement, then the 0.2 for Huysman.

Alexei Klimentov: Not an increase — a hire that was not made earlier. It was 0.4 with Yang and the fraction before he was hired, but from the beginning the Spradley replacement was planned at 0.5.

Verena Martinez Outschoorn: And did you put anything for storage R&D in the effort spreadsheet?

Alexei Klimentov: I put it as TBD — 0.25 for R&D, in FY27.

Verena Martinez Outschoorn: It's in there, you have it. It looks very reasonable — we're just asking so we can answer the questions.

Alexei Klimentov: Challenges. The reorganization went well — 95% of personnel decided to stay in SCDF, only two people preferred to remain in NPP, and one of them is now trying to come back. The main challenge is market trends. With the costs we know, we are mitigating: we delay procurement and extend hardware lifetime, and we redesigned equipment refresh plans to reduce upfront costs. Our situation today: we are not retiring CPUs bought in 2018; we are not retiring disks we planned to retire; and we are still waiting for a green light from management on the 6.4 petabytes of disk — we were asked to buy in two steps, we bought 10.6 PB but not the 6.4 PB. Our pledge is 49.5 PB. We also have a storage buffer in front of tape which is not counted as pledge — there's a long WLCG discussion about whether it should be — but we need it for high throughput to tape. With the 16 PB of disk due for retirement, we will not deliver the pledge for 2027.

Paolo Calafiura: Some basic questions on this table. The percentages — percent of what?

Alexei Klimentov: The age profile: each year's fraction of the total we have today. And one interesting number — the 60 KHS23 is about two-thirds of the total we purchased in 2018.

Paolo Calafiura: It may be good to add the pledge column next to the total, so you can show it visually — if it fits; otherwise say it verbally.

Ofer Rind (?): The mitigation discussed, given prices and budget stress, was not to retire

— the 2021 disk stays in production?

Alexei Klimentov: Of course we are keeping it. And if we retire it, we will not deliver the pledge.

Paolo Calafiura: On a previous slide you said SSDs went up 5x but HDD only ~30%. This disk is mostly HDD, I assume — the 6.4 PB purchase would be HDD?

Rob Gardner / Shawn McKee: It's always been our policy for bulk storage — we cost optimize with HDDs, and buy NVMe or SSD only for servers or specialized storage. The petabytes and petabytes of bulk storage are HDD.

Paolo Calafiura: So one more reason to focus next year's purchases on storage — HDD inflation is not insane. And looking at the table: the 24% block is 2021 disk, five years old — six by next year. We really should buy. Maybe split purchases to even out, but it sounds like we should at least try.

Verena Martinez Outschoorn: We should push to buy the 6.4.

Shawn McKee (?): The 10.6 was purchased, and the 6.4 was always part of a split purchase — planned for FY26, but we were asked to do it in two tranches. We need to find out whether Srimi has spent that money or whether it's still there.

Alexei Klimentov: Even so: if you order today, we receive in December.

Fred Luehring: Question on the storage table — the 67 PB projected for this year: does it include the 10.6 we purchased? The blue dashed box doesn't include it, but then the forecast of 67 can't include it either, because it would be 77.

Alexei Klimentov: The blue dashed box does not include the 10.6. You're right — with the 10.6 it was 77. We should fix that. Good point.

Alexei Klimentov: Next: many initiatives and ad-hoc projects this year — American Science Cloud (AmSC), Genesis, program-development work — requested more time from people. Even if on paper you have 100%, switching between projects takes time, and the increasing number of projects doesn't help. Services: Varnish is more for US ATLAS; CVMFS is for ATLAS and CERN. Now, budget. Shigeki and I think we will change the operations/capital split presentation next year, but this year we kept it as before. We managed to economize about \$442k on operations this year, but we have the deficit on capital for hardware. We can postpone it to future years, but it won't help much. Operations look good — we dramatically improved the situation, including the space-and-power discussion from last year.

Rui Wang (?): This power upgrade — is that something we contribute to, or is it folded into the power cost?

Alexei Klimentov: The data-center power upgrade is paid by the lab. We pay a fraction of power, calculated by the fraction of US ATLAS equipment in the data center — RHIC pays its fair fraction, Belle II pays its fair fraction, and so on. The question was always how it was calculated; we spent about a year getting the calculation right, and that's why it dropped by ~\$100k. Compared to FY24 it dropped even more — partly the calculation, partly bulk installation.

7

Paolo Calafiura: Small question on slide 30 — you use the words “FY27 guidance.” Guidance from Srimi, or...?

Alexei Klimentov: From Chuck. We asked Chuck, as before, what our budget will be, and we received a number. And then in August we were told there would be a \$1,048k cut.

(Brief interruption — Fred unmuted with a cat-related aside; asked to mute.)

Alexei Klimentov: Budget for hardware. The serious changes: hardware for general services and virtualization — we increased lifetime and moved it to FY27. Same for fabrics — we bought no hardware; moved to FY27. For storage, the famous \$640k we did not spend. Our forecast: to fulfill everything as in the past would be about \$2,128k.

Paolo Calafiura: What are the purple arrows?

Alexei Klimentov: We planned to buy virtualization hardware in FY26 for \$450k. We did not buy — we extended the warranty, which cost \$75k. Infrastructure equipment: the initial plan was to refresh 70% in FY26 (\$450k) and 30% in FY27 (\$200k). We didn't do it. In June we extended support for the 70%, and the current plan is to refresh 100% in FY27, which is \$650k.

Verena Martinez Outschoorn: Note that it's \$650k on the other slide and \$677k here — a \$27k difference.

Alexei Klimentov: If you want me to be precise to \$27k, I will do it.

Paolo Calafiura: And do we even know it still costs \$450? ... I'm stunned by these amounts, honestly. It's similar to what we buy for compute.

Rob Gardner: That all changed this year, Paolo. I'm not totally shocked — these are probably NVMe drives backing the VM system. It's expensive, premium hardware, but honestly it's probably not unreasonable.

Alexei Klimentov: And note we spent nothing in FY25, and from FY28 the virtualization hardware cost will be close to zero. I also have a comment that these are current prices, which can change. M&O activities: no big changes since last year. Milestones for FY26: we have delayed milestones, related to the personnel and hardware cuts — practically all related to storage. On working with 2.5: as soon as Kaushik and Zach have time to talk with Shawn and myself, we'll identify what we can do together. So far I don't know what 2.5 is going to do.

Shawn McKee: One possibility is to talk in person next week at Stony Brook.

Alexei Klimentov: Friday, Shawn and I are dedicating to a run-through of the 2.3 slides for all level-3s. In person at Stony Brook is best — it's not something we need to finalize before scrubbing; there will be milestone updates.

Kaushik De (?): Zach and I were also discussing this morning that we need to talk about possible FTE swapping or common work — so we have something to discuss next week.

Alexei Klimentov: New milestone presentation: high/medium/low priority (Shawn's convention: high = 1, medium = 3, low = 5), three pie charts — per sub-WBS, per category,

8

and FTEs per category — then a table summarizing counts and FTEs per group, milestone descriptions, and a link to the full milestone table. Risk analysis: the capital budget cut combined with current prices has two main impacts — huge stress on budget and pledge deployment, and increased operational load from running older equipment, which together with the personnel cut will impact operations. Data-center infrastructure has some issues by design. And the number of projects supported keeps increasing.

Paolo Calafiura: Back to the equipment slide — walk me through the storage plan again?

Alexei Klimentov: We don't have a plan, because we don't have a budget. I can't propose anything until you tell me my budget. Our initial plan was to buy storage; the current crisis intervened. The \$640k — maybe it's now \$700k. The early purchase, \$1,035k, was bought. The late purchase was postponed — \$640k was the end-of-May estimate.

Paolo Calafiura: And the FY28 number — \$2.2M for storage — is that replacing the 2023 bulk purchase?

Alexei Klimentov: Our bulk purchases were 2023 and 2024 — 22% and 21% of today's storage. On a five-year cycle, in FY28 we replace the 2023 tranche. The 2021 disk (24%) replacement is covered in FY26 by the 10.6 + 6.4.

Paolo Calafiura: Now I understand everything. So the big issue comes afterwards.

Alexei Klimentov: And we are running very old hardware, and we have started cannibalizing components from broken servers again.

Paolo Calafiura: Basic question: from Chuck you got FY27 guidance of \$3,965k, of which \$1,737k is capital, and your capital forecast is \$2,128k. And what are surplus and deficit on slide 30?

Alexei Klimentov: I want to emphasize we economized \$442k. Other experiments are allowed to move money between operations and capital; we never did that for ATLAS. We are absolutely transparent — we don't ask for more while planning to use operations money for hardware. But if US ATLAS allows it, it would help. Surplus: guidance for operations minus our operations forecast — \$442k. Deficit: the \$450k virtualization plus the deferred purchases — \$793k in FY26.

Paolo Calafiura: So the deficit isn't money you owe — it's the things you couldn't purchase. And the \$391k in the FY27 forecast corner?

Alexei Klimentov: The FY26 deficit, if we don't act, rolls into FY27, and there is also the FY27 deficit of \$391k — \$2,128k minus \$1,737k. The total accumulating deficit goes above \$1M.

Paolo Calafiura: And the operations economization — can you imagine getting it for FY27 as well?

Alexei Klimentov: We hope to stay at the same level, at least for FY27 — based not on electricity cost but on our fraction of hours decreasing.

Paolo Calafiura: I'll try to make a concrete proposal for reorganizing how this is presented. The crucial thing to emphasize: compared with Chuck's guidance and with what you had

9

after scrubbing, you economized \$442k last year and will do something of that order this year — and still you are in the hole on hardware big time.

Alexei Klimentov: It is a \$1M cut — plus about \$1.2M in the two previous years.

Paolo Calafiura: We need to make very clear that you basically survived a couple of really savage cuts, that there is a debt that has been accumulating of things you couldn't do, and that there is nothing left to cut, essentially. I see a lot of green and I didn't appreciate that you were \$1M in the hole, depending how you count.

Shawn McKee: A couple of closing things. We have Friday afternoon, 1:00 to 3:30 Eastern, targeted as another pass — at least Alexei and me; others welcome, let us know. It's to make sure everything is updated based on the pre-scrubbing responses, and we'll go through all the level-3 areas.

Alexei Klimentov: Could I present my final set of slides on Friday too?

Shawn McKee: Yes. Second: Rob provided some less-detailed 2.3.4 slides this morning for consideration — maybe Wei and Rob and others should look and see whether we can converge on something with less detail but all the needed information. Third: there's a problem I identified with 2.3.5.4 cybersecurity — it turns out that in the next CA we do not have any new money for that effort, and what appears to have happened is that 0.2 FTE was taken from elsewhere in 2.3.5. We have to think about what this means and what we're going to do about it.

Verena Martinez Outschoorn: And please — everybody take another double, triple look at the effort spreadsheet and make sure it matches, and assume nothing moves to 2.5, because we're trying to roll up the budget for Srini and we're already behind. If it's not in the spreadsheet, it's very difficult for us.

Shawn McKee: Verena, I sent you an email asking about Ilija's effort — the 2.4.3.2

question. **Verena Martinez Outschoorn:** I just saw it — I'll respond.

Rui Wang (?): Shawn, I noticed you were updating the spreadsheet to reflect effort after moving it to 2.5. I was revising my slides to show the full effort, as Kaushik suggested — which are we doing?

Shawn McKee: I put comments on each change in the 2.3.5 FY27 rows — those should probably be reverted, then.

Verena Martinez Outschoorn: I think that's safest — otherwise people fall through the cracks and numbers go missing. Everyone should present their FY27 proposal with the full effort side-by-side with FY26, showing the proposed move to 2.5, and we do the actual move after.

Ofer Rind: While I have the floor — Alexei, on the data-center power instability slide (27): is it worth showing the monitoring plot with the low job numbers during that period? We provided our pledge, but we were down thirty percent or so during the instability. It shows a real effect on the computing we provided.

10

Qiulan Huang (?): A question for Alexei: for 2.3.4 BNL personnel, do we currently have four names? Who should I ask for the final names and percentages?

Alexei Klimentov: You have Tom Smith, you have Shalom(?), and you have Ofer. If there's a fourth name — Louis — that's news to me; we never discussed it. Discuss with Ofer, since Louis is in his group, but we need a consistent picture for scrubbing.

Ofer Rind: I never discussed having him work in that group. He was working on containerization, partly for SDCC but mainly for the traineeship Verena is running. I don't know how that became a 2.3.4 milestone — 2.3.4 will benefit from the effort, but it was never meant

as a 2.3.4 deliverable. I'll follow up on the details.

Alexei Klimentov: Can we have a dedicated discussion about 2.3.4 on Friday? And I hope you'll have the utilization plots that Ofer and Paolo asked for.

Shawn McKee: And again, I'd encourage everyone to look at the other slides Rob put together. We don't have much time. Shawn, Alexei, and Paolo — can you stay on for two more minutes?

All: Thanks, everybody.

Sources and Methods

Inputs: (1) A Granola transcript of the July 8, 2026 US ATLAS Facility Coordination call, labeled only "Me" (Rob Gardner) and "Them" (all other speakers); (2) the attached PDF of Alexei Klimentov's WBS 2.3.1 scrubbing slides (dated July 14, 2026); (3) WBS 2.3 project context. No Zoom AI summary was provided.

Method: The transcript was cleaned against the slides, which served as the authoritative reference for names, slide numbers, and figures. Consecutive utterances were merged, phonetic errors corrected, and broken sentences completed where meaning was recoverable. The summary was written from the cleaned transcript.

Significant corrections and inferences:

- Speaker attribution is inferred, not given: Rob indicated Alexei presented and Paolo and Verena asked most questions, with some from Shawn and Ofer. "Paolo" is taken to be Paolo Calafiura and "Verena" to be Verena Martinez Outschoorn (US ATLAS management, consistent with "Paolo and I will get grilled" and rolling up the budget for Srinu). Blocks marked (?) are lower-confidence guesses (e.g., the 2.3.4-names question is tentatively attributed to Qiulan Huang; the power-upgrade contribution question and the 2.5 slides question tentatively to Rui Wang).
- Name corrections from slides: "Naved/unaware/Armware/on their/under" -> N. Anwer; "Spradly" -> Spradley; "Alex Suntros and Zayo Young" -> A. Undrus and Z. Yang; "Stefan Hismom / Steven Heusman" -> S. Huysman; "Shigechic / Shiggy" -> Shigeki (Misawa); "Elia" -> Ilija (Vukotic, presumed); "Genesys" -> Genesis; "Sean" -> Shawn; "Virena/Veronica" -> Verena; "Alfred" -> Fred. "Shalom" in the 2.3.4 BNL names exchange could not be resolved and is flagged.
- Figure corrections from slides: "3,960,000" -> \$3,965k FY27 guidance; "1.7" -> \$1,737k capital guidance; "2,100,000" -> \$2,128k capital forecast; "1.28 minus 1.07" -> \$2,128k

11

- \$1,737k = \$391k; "\$4.50/\$2.27/\$6.77" -> \$450k / \$227k / \$677k; slide numbers realigned to the PDF (FTE mismatch discussion is slide 22; budget table slide 30; data center slide 27).
- Garbled transcript rendered "two four three point two" (kept as 2.4.3.2) and "February" in the Louis exchange, which context indicates meant "2.3.4"; corrected.
- Fred's unmuted cat interlude was compressed to a parenthetical.
- "Console from management" was interpreted as awaiting a green light from management on the 6.4 PB purchase.

