Question Count and Status

Subcommittee		Answered in BO	Answered in Q&A
PM	5		1-5
C&S	8	All - offer to review for full committee	
CE	10		
OT	9	4	
MTD	5		
TD	2*	1,2	
Total	35+2*		

Project Management

1. Please provide a summary of costed and un-costed resources that need to be filled by hiring now and in the future and when are they needed?

The Costed and Uncosted Labor resource profiles per institute are shown below:

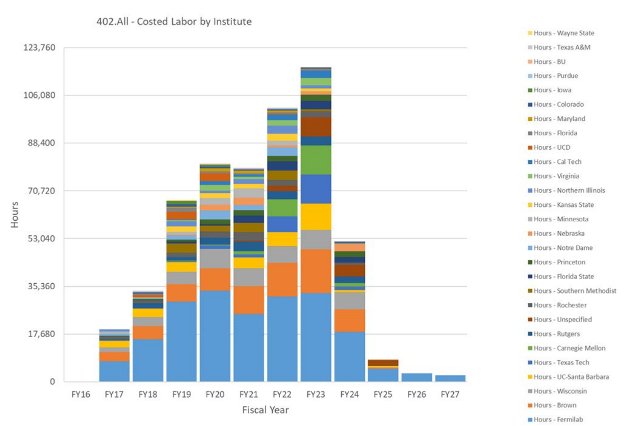

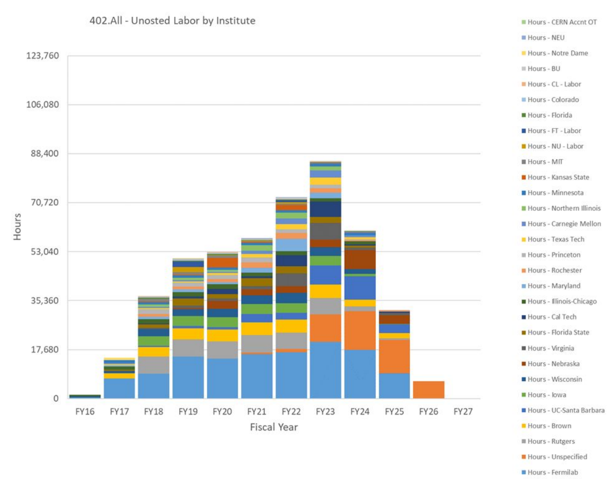

The summary is that for both categories, the labor needs build through prototyping to a plateau in FY19-21 before entering production in FY22-24. For the costed labor, currently for the scope of work we are doing, we are fully staffed, although in production several assembly sites (Brown, Fermilab, Nebraska, TTU, CMU) will need to add technicians to their teams to handle production. These sites have already at least one full time technical staff involved in the prototyping and development of procedures, building up expertise to become the “site guru” who can also train new technical personnel. In addition to technical staff, university sites envisage hiring a small cadre of undergraduates to supplement uncosted labor in QC activities.

For the uncosted labor, the same trend appears, peaking at about 50 FTE total in FY23, roughly 10 from Fermilab and the rest from the institutes, in FY23 during the height of production. For scale, as part of last year’s portfolio review, U.S. CMS tabulated the scientific labor pool available to work in FY17, which included about 55 FTE of Fermilab uncosted labor and 327 FTE of uncosted labor from the universities, for research, operations, and upgrades. Even if this pool is cut by 10%, if on average a U.S. CMS person works only 15% on the upgrade, the uncosted labor needs are met in the peak year. For hiring, the turn-over process is continuous on the typical 3-4 year timescale of a student or postdoc, but out of phase at the institutes, such that gaps at any particular institute do not create a large perturbation to the project.

2. What are the risks associated with the MTD performance? The risk registry only addresses delivery and component issues.

The challenge of the MTD is not so much that it is an unprecedented detector, but rather it has unprecedented precision. This precision is driven by the component performance, and the physics studies show significant unique physics gain even if the precision only reaches the 60-70 picosecond level, which we believe is well within the capabilities of the current design.
On the U.S. scope, the critical items in terms of MTD performance are the ASIC and the SiPM. We already have risks on the performance specifications of the ASICs, which add an additional round if the performance is not satisfactory, for both prototyping (RT-402-8-01-D ETL - Additional FE ASIC prototype cycle is required) and preproduction (RT-402-8-03-D ETL - FE ASIC does not meet specs - needs another pre-prod run). Similarly, there is a risk for SiPMs (RT-402-8-14-D BTL - Problems with SiPM vendor) which addresses substandard SiPM performance. Finally there is also risk for underperforming LYSO (RT-402-8-04-D BTL - LYSO matrices not meeting specifications) that affects the BTL Assembly.

3. What are the scope contingency items, the priority and the impact of each of them? What would be next on the list if you need another $5M?

Without knowing the actual situation, it is not possible to try to pin down what exactly the optimal descope solution is, thus the project has a list of options. The principles behind the choices of possibilities for scope contingency are

1. Leverage the broad collaboration by choosing items already being done elsewhere, or that are straightforward commodity procurements, such that the lost scope can be easily adopted by a collaborator.
2. Have a variety of options to choose from multiple L2 areas
3. Make decision points as late as possible
4. Attempt to preserve U.S. Core contributions as much as possible

Currently the 5.1M consists of the following items, each marked by a “Downscope Milestone”:

- Travel: $216,613 -- project personnel reduce travel to CERN, risk of decoherence with international effort
- Outer Tracker Module Assembly Batch 10: $649,448 International partner would pick up assembly, or Tracker would be missing ~3% coverage.

- CE Testbeam wedge: $575,625 Removes capability to study performance in testbeam to understand detector behaviour.
- CE Commodity Power supply System $2,034,762 International partner procures power supply, or several layers are not powered
- TD: DAQ procurement $1,409,293 International partner procures Data Storage, or CMS does not record data
- TL: Last 25% of BTL Assembly $132,631 International partner would pick up assembly, or BTL would be missing ~15% coverage
- TL: Last 25% of ETL Assembly $92,331 International partner would pick up assembly, or ETL would be missing ~12% coverage

There is no prioritization at the moment other than perhaps the travel, because as the HL LHC Detector Upgrade project is part of a larger international upgrade of CMS, decisions on scope reduction cannot be done unilaterally, but have to be carried out at the international level, where the flexibility of having additional collaborators can be fully exploited to come to the best conclusion as to how to move forward and preserve performance. The archetypal document for such discussions is the CMS Phase II Upgrade Scope Document (CERN-LHC-2015-19) , which discusses modifications to the Technical Proposal (CERN-LHCC-2015-010) under reduced funding scenarios. While because in the interim the design has progressed well beyond the Technical Proposal, some of the options proposed in the Scope Document are no longer feasible, the same process coordinated through Upgrade Coordination would take place on a smaller scale when considering the priorities among different options for down scope. This strategy has already been approved by the Upgrade Coordinator.

For more scope reduction, there are several possibilities we can imagine, but the question is at this point “which limb do you want cut off?”, and the full impacts in terms of cost and physics performance are not quantifiable overnight

- Further cutting assembly labor in OT and TL, putting additional burdens on the international collaboration or increasing the unrecoverable loss of acceptance
- Cut assembly of odd-size CE-E modules, with unrecoverable loss of acceptance
- Choose less performant FPGAs in the trigger, reduce both L1 rate and “physics purity” of L1 data selection
- Completely remove construction spares and teststands from all L2 areas. This would result in reduced detector performance since we would be forced to accept potentially substandard components on the detector, or introduce holes in the acceptance
- Stage the ETL construction, since it can be added later (but as most of the cost is in the ETROC engineering, this is not a big cost saver)
- Remove one or more layers of the CE, unrecoverable without international adoption
- Remove the Correlator Trigger, which can be recovered later, with loss of physics performance
- Remove the Outer Layer of the Flat Barrel, unrecoverable

Finally, at this point,the international collaboration would also start to question the necessity of preserving $10M for installation and commissioning when the U.S. is failing at fulfilling the MOU responsibilities and damaging its reputation for delivering on scope, so in addition this would need to be discussed between the DOE, the project, international CMS, and CERN.

4. What are Lucas Taylor's additional responsibilities and how much time does he currently spend on CMS? Is there an opportunity to train a Risk Manager?

Lucas Taylor's primary role is **HL-LHC CMS Upgrades Associate Project Manager**, for both NSF and DOE scope, which takes approximately 70% of his time. The remaining 30% is spent on his duties as **Fermilab Risk Manager** and **PIP-II Risk Manager**. On CMS, he is responsible for the development of the Cost, Schedule, and Risk aspects of the project, working closely with PCS and PMs.
The **Fermilab Risk Manager** job entails the support and development of Lab-standard risk processes and tools (including the web-based risk register tool and the Lab's Strategic Planning tools), MC modelling, and risk reviews, workshops and mentoring. It entails monthly Fermilab Risk Management Board meetings and occasional retreats or workshops with senior Lab management.
The work as **PIP II Risk Manager**, is more substantial but intermittent, with peaks as PIP-II approaches reviews or major workshops and lulls between. Lucas and OPSS Management are working to broaden the base of people who can work in risk management using Lab-standard methodologies and tools.

5. Could you make a table with for each L2 the big external dependencies, eg sensors, lpGBT, LYSOs, MPA chips, etc, when they are needed, when they are expected to be available, and the
distance to the critical path?

We already have such a list, we currently call it the "watchlist" but intend to rename it because that word has a different connotation for baselined 413.3b projects. It is a live google doc but we do take a snapshot of it, and store that at https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13742 which is also available from the review documentation page under the "Change Control and Reporting".
In addition, we have about 275 External Milestones to keep track of externals (see Lucas plenary talk, slide 17), although not all of them are used in the logic flow.

Cost and Schedule

1. Critical Path scheduled from P6 by L2 showing Total Float.

The L2 areas in P6 are completely independent without logic links between them. The project has a critical path for each L2 area as shown in the plenary presentations. The meaningful critical paths for each L2 area terminate on the delivery of the threshold KPP scope (with well-defined deliverables). The objective KPP scope includes integration and commissioning tasks, which, by their very nature, will make use of all time available before the installation of the detectors into CMS.

To determine the critical paths, we constrain each threshold KPP milestone to its technical finish date so that total floats for each L2 area are calculated with respect to the completion of the threshold KPP. Hence the project critical path comprises all activities with Total float = 0d.

The resulting Gantt chart, with columns showing Total Float and Base Cost (i.e. Planned Cost) is here: https://cms-docdb.cern.ch/cgi-bin/DocDB/RetrieveFile?docid=13245&filename=2019-10-22---HL-LHC-CMS-Upgrades---Gantt-Chart---with-Total-Float-and-Cost.pdf&version=30

Notes on changes to P6 since the frozen versions that were used to generate review documents:

(1) *A trivial bug was fixed in 402.4 CE. This related to a constrained milestone logic and lag that resulted in spurious negative float appearing for some activities in 402.4 CE. This bug did not affect any costs or dates.*
(2) *An unrelated minor schedule fix that was made after the cost documents were frozen for the review advanced two activities in Calorimeter Endcap, resulting in a slight change in cost (reduction of $6k due to less escalation) in the latest P6 files compared to the frozen cost book.*

The xer file after these trivial fixes is available here:
https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13554

2. Provide Need by dates and Requested dates for each L2

Only MTD has a “request-by” date, because of the multiple options for integration into the Tracker Support Tube open to them. We use the need by date to use to measure float.
The following graphic shows, for each threshold KPP milestone, the corresponding CMS Need by date.
The yellow bars indicate schedule contingency between the technical early finish and the CMS need by date.

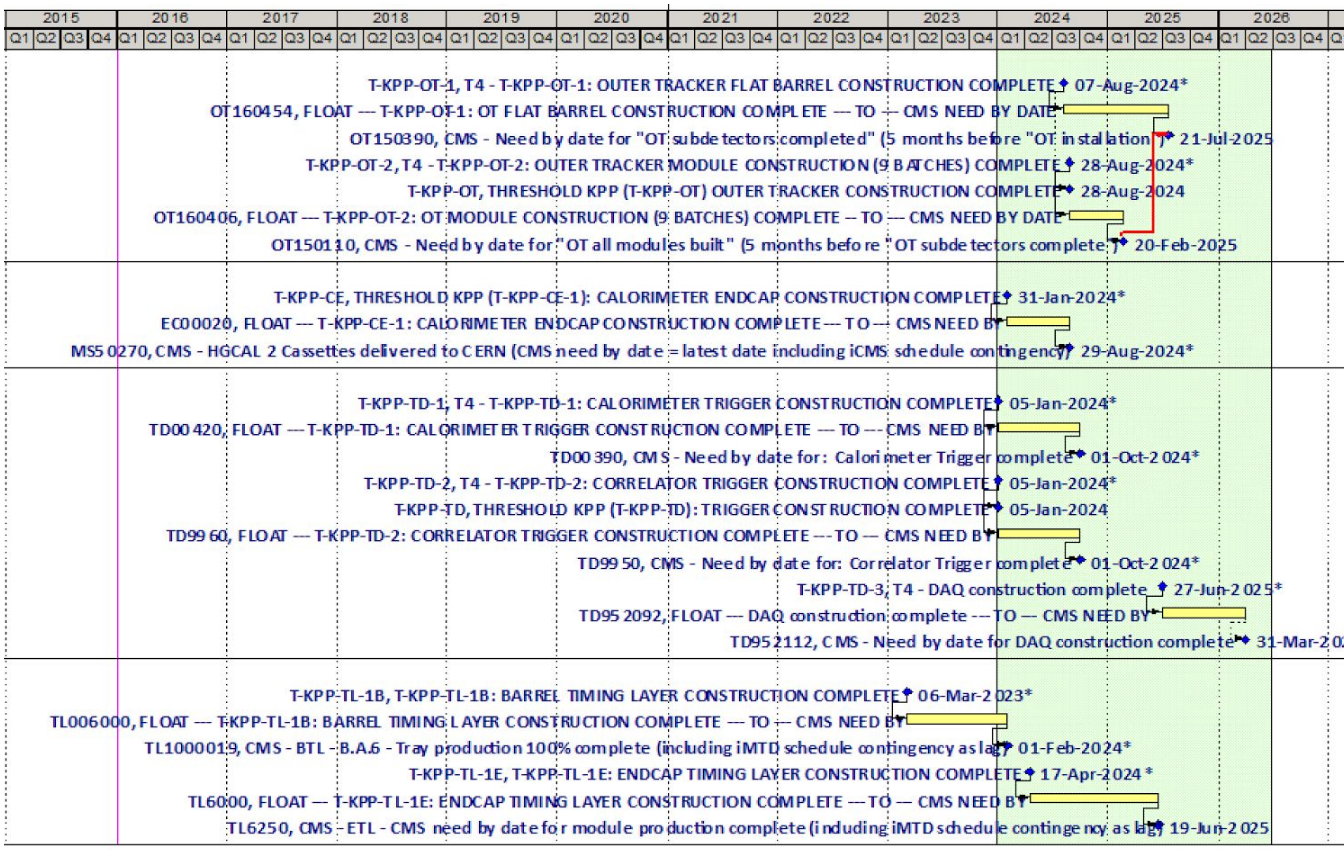

3. High -Low cost Range excluding scope contingency.

This is ongoing but suffered from an unplanned mandatory operating system upgrade which hampered P6 access (other homework was enabled via a phone hotspot).

4. Fuse Analysis

The output of the Acumen FUSE analysis is available here:
https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13917

5. How are you managing staff availability

Staff availability is managed differently depending on where the staff is located. At the Institutes, the PIs sign a SOW that states both what the deliverables are and what resources, split by M&S and labor, are provided. It is the responsibility of the PI to then hire the staff to achieve that scope. Should they be unable to do so, a discussion happens with Project Management to see how to resolve the problem, typically by moving the scope of work elsewhere. At Fermilab, staff is heavily matrixed, so staff availability is negotiated between the Division resource manager, the Department heads, and the L2 or L3 managing that part of the program. The “Roles” field in P6 is used for all Fermilab labor, to identify either the person or the type of person (tech, engineer) needed for the activity, and this is collected on an annual basis to establish the project’s resource request to the Divisions. When conflicts arise, the L2/L3s discuss with the Project management and the Chief Project Officer, who has the role of Labwide resource arbiter, to explore solutions. Most recently, the Chief Project Officer established the ASIC PMG, a monthly meeting where the resource usage and upcoming demands of the ASIC engineering team is reviewed and discussed between the various stakeholders, and if necessary, prioritized.

6. What is your plan for managing contributed /uncosted labor?

Contributed labor is managed by tracking milestones for the deliverables they are involved in creating. For instance, there are milestones in the trigger for versions of software emulation and algorithm development, which are done chiefly by contributed labor. Another example from Phase 1 was Module production, where the QC was done by contributed labor led by former UIC postdoc Doug Berry (now a Fermilab scientist and L4 in OT), and the progress of tested modules was followed weekly:

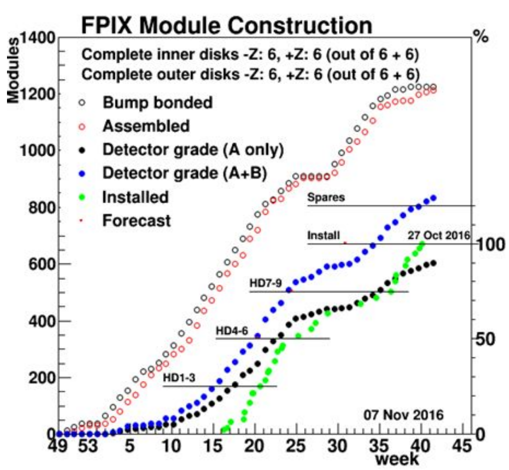

FPIX module production - the step between assembled and graded was dominated by contributed labor.

As a final example, this is a photo of “Nadja’s Ninjas”, the undergraduate team lead by former Fermilab post-doc Nadja Strobbe (now a L4 in CE) that did the QC testing of the ~ 800 QIE Cards for the Phase 1 HCAL upgrade, completed in about 6 weeks

We are confident that this method of tracking the progress of uncosted labor ensures that we will see the delays associated with uncosted labor shortages should they occur.

7. Provide gantt chart with total cost and float added for project.

See answer to question 1 regarding critical paths for L2 threshold KPPs. The corresponding Gantt chart with columns showing Total Float and Base Cost (i.e. Planned Cost) is here: https://cms-docdb.cern.ch/cgi-bin/DocDB/RetrieveFile?docid=13245&filename=2019-10-22---HL-LHC-CMS-Upgrades---Gantt-Chart---with-Total-Float-and-Cost.pdf&version=30

8. Breakout of Total Budgeted Cost by element of cost (base cost) overhead and escalation by Level (2?) by FY

Institute labor rates are collected from all the US-CMS PIs. The base cost, the fringe costs, and the applicable institute F&A rates are also specified as distinct items. These data are all available here: https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13284

These rates are then imported into Cobra by Fermilabs Office of Project Support Services, which becomes the definitive source of the labor rates applied to university personnel. However, the Fermilab Cobra system does not maintain -- as distinct data -- the institute F&A rates and the institute fringe rates; instead these are multiplied together to obtain a fully-burdened rate for each flavor of institute labor.

Cobra then applies Fermilab overhead (pass through) and escalation to the fully-burdened institute labor rates in the same way as for Fermilab resources.
The project resources -- allowing for indirect costs and escalation -- are also available in P6 (this is currently the default cost engine for the project, although Cobra will be the ultimate source as we move to EVMS). These resources in P6 are obtained by exporting the corresponding data from Cobra and importing them into P6. Therefore, P6 also has no knowledge of the details of the institute indirect and fringe costs (just the fully-burdened net cost per resource). It is therefore not possible to show the base labor rates separated from overheads and fringe for institutes in P6 or Cobra.

The full details of all the input data (institute labor base rates, indirect costs, and fringe) are however available here: https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13284

For other resources (M&S and Fermilab labor), Cobra does manage the various cost elements (indirects and escalation) as distinct data (i.e. not combined into fully burdened and escalated numbers). At this time, we are using P6 as the default cost engine. We can drill down on individual P6 rates to show the committee specific resource rates including overheads, as a function of fiscal year (hence showing escalation). We have requested the P6 support team to export all P6 rates, as a function of each fiscal year. This is not a custom DB export and may therefore not be available on the desired (short) timescale.

Outer Tracker

1. Quantify the impact of automatising the module assembly in terms of time and labor costs - labor costs should also be detailed in terms of technical / generic. As this one question from last year review which was requested, it would be useful to have quantified information which assessed the current situation.

From the Risk Registry (RO-402-2-03-D): from the Module Assembly Workflow_v7-June2019.pdf document the fraction of labor hours for 2S/PS assembly that could be saved by implementing metrology measurements and hybrid gluing tasks on a gantry robot is about ⅓. Since the P6 estimate for 10 2S and 10 PS production assembly batches is $2.4M (fully burdened), this would suggest a savings of $800k in labor. However, technician labor at Brown would be replaced with undergraduate labor, which would incur some cost, so the total savings would be reduced to about $700k (for both FNAL and Brown). From this the cost of a gantry system, laser, and other related equipment for Brown should be subtracted. (All of the equipment is currently in place at FNAL.) This is approximately $200k, so the total savings due to gantry automation is about $500k.

2. Quantify the impact of having or not an irradiating facility or getting it with a 6-12 months delay. Quantify that in terms of costs and time and add mitigation or fall-back plans.

In the highly unlikely case that we don't have the ITA at Fermilab, or that it is delayed until after the start of sensor production, we will still do neutron irradiation at the Rhode Island reactor. Neutron irradiation is more harmful to sensors, so it will provide a conservative estimate. For proton irradiation during production, we will be forced to send our test structures to Karlsruhe, Germany. In addition to having to pay for shipping and irradiation in Karlsruhe, we would incur a delay of a few weeks, so if there is a problem with one of the batches, the feedback will still come during the duration of this batch, but our reaction time to potential problems would be delayed by a few weeks with respect to having the ITA at Fermilab. In that case we risk a lower yield for a longer time. This is covered by our risk about a lower sensor yield during production (RT-402-2-01-D) with a probability*impact of $48k and 1.8 months.

3. Clarify how the various module components are going to be tested in terms of thermo-cycling, irradiation, and batch variations not only for sensors [international basis?]

All ASICs will be radiation tested. For ASICs there is only one production run, no batch-by-batch variations. Every single hybrid will undergo extensive long-term testing with multiple temperature cycles at the international testing sites (CERN & INFN). The module mechanics materials, such as CF, glues, kapton, will be qualified for radiation hardness.

Within the Modules subproject there are plans for reception testing FE hybrids and service hybrids to ensure they were not damaged in shipping. Mechanical components will also be inspected to make sure they are in spec. There are no plans to reception test all sensors (other than the PS-p testing that is implicit in MaPSA probing) as any benefit in flagging issues would be far outweighed by the risk in the handling of the sensors.

4. Clarify MaPSA situation and risk factor - is there a mitigation or fall-back plan in case all vendors fails or go below expectations? Add this to the dedicated presentation.

We believe there is only a very small probability that all vendors we are vetting will fail to deliver an acceptable product. Technically, the bump bonding of the MaPSA is less challenging than previous projects we have experience with, such as the original FPIX and Phase1 FPIX modules. The UBM and bump placement R&D is done and these vendors are secure (HPK and StatsChipPak). The flip chip assembly is not technically challenging, but the form factor is unusual for most vendors. Initial dummy parts we received from the first vendor show good quality of the bumps themselves in terms of uniformity of size and placement. The actual assembly did not cause the delay in the last round. This was caused by the vendor working through problems serially (picking parts off dicing tape, slow delivery of waffle packs, communication with the dicing vendor). The actual module assembly went smoothly and took only a few days. HPK assembly is delayed because they were waiting for parts from the first vendor.

We have had acceptable setup parts from two vendors. We are now starting to test the first fully active modules. An outstanding issue is the transfer of MaPSA probe testing from Fermilab to a commercial testing vendor, which had to wait until active parts were available.

Additional risk mitigation can consist of probe testing of the bump bonded modules at the vendor, allowing for replacements of individual, faulty chips.

If the best vendors do deliver acceptable parts, but at a slightly lower quality than our specifications for example with respect to fraction of functional pixels, we would have to evaluate the impact on physics performance through simulation. However, even having a factor of two more bad bumps than specified for example, will have only a minimal, if not negligible, impact on tracking efficiency and resolution. We do have a risk that covers lower performance of MaPSA (RT-402-2-09-D), with a probability*impact of $76k and 0 months.

5. Skilled personnel - do you have a way to quantify the risks for loss of key-skilled personnel / difficulties in recruitment in the next few years - particularly related to schedule delays

We have a risk that covers this (RT-402-2-90-D) with a probability*impact of $73k and 0.3 months. There are only a very small number of very specialized technicians, namely for micro-assembly of the modules and wire bonding. They will be harder to recruit and involve longer training times than regular technical personnel. Personnel can be cross-trained, exchanged, and lack of personnel at one site can temporarily be relieved by an increased throughput at the other site.

6. Hybrid delays - impact of this should be better explained. Similarly for the sensors. One question asked to CE which is also relevant for OT: what is the relative schedule impact of the schedule drivers. For example, set sensors to 0 and assess the others relative to that

This is captured by risk RU-402-2-01-D OT - Uncertain performance of Hybrids vendor. For hybrids (which encompass lpGBTs, some external ASICs, and DCDC converters), a delay of the start of delivery of the latest available type of hybrid up to 84 working days does not affect the Threshold KPP date.

7. Can you specify which ‘small amount of R&D’ is effectively remaining?

No real R&D is left, except for a few design development tasks and value engineering:

- Sensors: none
- Electronics: commercialization of MaPSA testing, MaPSA underfill choice
- Modules: AlN, adhesive kapton, automation (value engineering)
- FB: final dimensions of tilted flat barrel interface, cable and fiber routing

8. LPGBT: there was a mention of possible delays, what is the exact impact of that on the overall schedule ?

We receive hybrids that are fully populated with all the discrete components, including lpGBTs from CERN. A delay in their delivery is covered in our hybrid delay risk (RT-402-2-01-D) with a probability*impact of $272k and 4.7 months. For the delay, see answer to 6).

9. What is the status of the simulation of the tracker in terms of material budget in general, and specifically in terms of cabling, services, infrastructure ? Are there test measurements done on a regular basis also of uncritical material like cables? Any study on this made in the past could be discussed tomorrow.

Detailed material is already included in HL-LHC GEANT-based simulations. Our experience from the Phase 1 FPIX upgrade for example, shows that final, actual components, including screw and all cables will be weighed and recorded during production of the final detector components and then will be used to refine the already existing material budget in the simulation. In addition, we continuously verify the material in the simulation with data-driven material budget maps, for example for Run 1 and Run 2. It has to be noted, that this is not really part of the US scope of the Outer Tracker.

Calorimeter Endcap

1. What are the explicit threshold and objective KPPs (i.e. deliverables). Are these specified in the WBS dictionary? Clarify slide 14 in plenary presentation?

See table. These numbers are well known and understood within the relevant L3 areas. The project asked specific numbers of e.g. modules not to be included in the KPPs, or the WBS dictionary before baselining.

2. What fraction of the R&D remains to be done?

The project is currently in the late stage of R&D and most areas are moving into the stage of detailed engineering or process optimization. The aspects which are currently completing the R&D phase include the SiPM packaging and the tooling for scintillator wrapping. Areas of process optimization include final tuning of the 8” sensor line and tile module pick-and-place optimization. Major R&D which is complete is the SiPM device structure, the sensor design, and identification of appropriate scintillator and wrapping materials. Of the R&D required since the conception of the project, only a few percent remain.

3. Describe what was given up in the simplification of the ECON, its impact on performance and how the agreement with international CMS was orchestrated.

First, we mention that the simplification of the ECON also allows a major simplification of the system design through consolidation of complex circuitry into a single, small board (the motherboard engine) leaving the simple motherboard wagon as the only components with a relatively large number of variants.

The primary disadvantage of the new ECON design is that 15% more 10Gbps optical links are required to move the trigger data off the detector within the ECON latency budget because the updated ECON allows sharing of bandwidth over a smaller area of the detector (1 module vs. 3 modules). It was the judgment of CMS that the advantage of the reduced complexity+risk for the ECON and the overall electronics system outweigh the disadvantage of this (easily understood) increase in link count.

The agreement was orchestrated through a series of internal and external reviews. The project received a recommendation at the March 2019 CD-1 Director’s Review to consider a simplified ECON (without 10Gbps transmitters). At the same time, CMS informally requested we investigate ways to simplify ECON. At a May 2019 ECON review (organized by CMS), the US team presented a unified plan for the current simplified ECON/system architecture. Additional studies of the impact of the new architecture on trigger data latency, total link counts, and ECON trigger algorithms were reviewed at an HGCAL workshop in June 2019. The new architecture was adopted by CMS in early July 2019.

4. What people and expertise is available to use the software verification tools which are planned to be used to allow the ECON chip to be fabricated with a single iteration

Verification will be performed with the industry standard Universal Verification Methodology (UVM) procedures and tools. In Summer 2019, the Fermilab ASIC group brought in a consulting firm to provide a weeklong UVM training and evaluation of Fermilab’s existing UVM methods and tools.

The Fermilab ASIC group’s primary UVM expert is an ASIC engineer with 20+ years of digital design experience. He is working on ECON verification full time. He is also coordinating the work of a second ASIC engineer (50% on ECON) and a firmware engineer (60% on ECON), who are both less experienced with verification but ramping up quickly, as expected for experienced engineers. A post doc physicist is providing input and validation UVM “sequences” from physics simulation and python-based emulation of certain ECON blocks.

It is expected that this group of four will complete ECON verification on schedule, but we are looking to recruit a software engineer from Fermilab’s Scientific Computing Division to bolster Fermilab’s UVM capabilities, which would allow additional flexibility for ECON.

5. Give the specific example of the chip which was designed fabricated more or less successfully in a single iteration and characterizing the similarity between it and the ECON

The ColDATA ASIC was recently designed/fabricated and found to be effectively error free in a single iteration. ColDATA and ECON ASICs are both ~100% digital ASICs, which use the same process (65nm CMOS), designed by the same team using the same Cadence tools and digital design flow and verified with the same UVM tools and procedures (though some improvements will be used for ECON). In addition the functionality is similar: ColDATA receives data from the front-end ColdADC ASIC (similar to HGCROC); formats, encodes, and serializes the data; then transmits the data on two 1.28 Gbps differential transmitters (similar to the ECONs 12x 1.28 Gbps elinks). Both ColDATA and ECON slow control uses i2c and fast control is achieved with custom command interpreters.

6. Describe the engineering resources (i.e. people and competition with DUNE for them, also noting that there is no competition with the ETROC) needed to complete the ECON.

The required ECON engineering resources are shown on slide 27 of the ECON breakout session presentation. The design team is composed of 3 Fermilab engineers + 1 visiting engineer + assistance from European collaborators (Split and LLR, Paris) who will provide confirmed-synthesizable verilog for the ECON "best choice" algorithm. The total effort rate for design and verification (not counting European collaborators) is approximately 4 FTE.

Because ETROC is a primarily analog chip and ECON is primarily digital, the competition for resources is low. The primary competition is between ECON and DUNE who both need engineers with expertise in "implementation," which is the conversion of RTL into physical circuit design using digital design tools. ECON currently has sufficient implementation resources, but we anticipate that DUNE will require additional implementation resources in mid-2020. If ECON is delayed, there will be modest conflict for those resources in mid-2020. Fermilab is addressing this primarily by hiring 2-3 engineers, 1-2 with implementation experience, while looking into training other engineers already in the Fermilab group.

7. The project has had some setbacks with obtaining the needed resources, notably ASIC engineering. What steps has the project taken to ensure that this will not happen in the future and that resources will be available for the duration?

The general approach is to work with group leaders at each of the collaborating institutions to identify the resources available relative to those required and to identify possible future conflicts for those resources or other reasons why they might not be available in the future. We try to anticipate our needs before they are actually required and work with group leaders to ensure that they know what we need and make plans so that they can be made available when we need them. Some resource conflicts can be addressed by rescheduling work that is sufficiently far from the critical path for the work involved to make such a strategy possible. Resource commitments are documented in annual SOWs with each institution, the delivery of these resources are monitored by the appropriate L3 or L4 manager, and actions are taken if commitments are not met. For key resources, i.e. people with particular expertise on which the project depends, several strategies are followed: possible backups are identified in advance to the extent possible, e.g. identifying other groups who may have similar expertise; we may encourage the relevant institution making new hire(s) in advance of there being a shortage; or others within the HGCAL group could work to develop relevant expertise.
Two key examples of key resources on which we rely are ASIC engineers and the SiPM expertise in the Notre Dame group. Issues with the Fermilab ASIC group and how they are being addressed have been discussed above. The Notre Dame group is responsible for both the SiPMs in CE and BTL. We are aware of the issue and have taken care to schedule the relevant tasks such that the group is not over-burdened at any time in the current schedule. This process is facilitated by the fact that M. Wayne is the responsible L4 for SiPMs in both areas and will be following up in case of any adjustments in the schedule of either project.

8. Provide the date at which it which it will be known whether the expectation of a single iteration of the ECON will be successful and how you might accommodate/mitigate the resulting delay.

The ECON prototype will be submitted on 08-MAY-2020, and ECON prototype testing (bench, radiation, and system) will occur 21-SEP-2020 through 15-FEB-2021. The ECON production will start on 20-MAY-2021. Thus, it is most likely we would learn that a second prototype submission is required sometime between 21-SEP-2020 and 15-FEB-2021. The delay associated with the fix and resubmission is 6-9 months depending on the time required for diagnosis and modification of the design.

There are a number of ways to mitigate a possible delay in the ECON: We are working on a version of the design of the motherboard with a separate concentrator board (separate from the engine and wagon) that can be completed later; We will make a board like the concentrator board for testing the chain with the same electrical and mechanical connections; We can assemble cassettes without the concentrator board and complete as much of the testing as we can and stack these cassettes to wait for the concentrator boards to complete; We can speed up assembly and speed up the final testing of completed cassettes; For some subset of the final cassettes we can ship the cassette components directly to CERN for final assembly and testing there.

9. How long will it take for a second ECON iteration - we heard ~9 months but it would be good to have it written down.

A second ECON prototype iteration would require 6-9 months depending on the time required for diagnosis and modification of the design. The fabrication and packaging requires 4 months, and the diagnosis and design adjustment requires 2-5 months.

10. What is the relative schedule impact of the schedule drivers. Set sensors to 0 and assess the others relative to that:

ECON -- There are 4.5 months of float between ECON delivery and the point where motherboard delivery starts to affect cassette delivery.

lpGBT -- The lpGBT current schedule has two deliveries of lpGBTs. The first large-quantity delivery is scheduled for Q4 2021, which is the same time frame as the ECON delivery. The second batch is delivered in Q3 2022. If the CE project receives at least 3000 units from the first delivery, there will be no schedule impact on final delivery of cassettes. A delay of up to 4.5 months on the first lpGBT batch can be absorbed in the project without a delay on the final delivery of cassettes. If the CE does not receive lpGBTs until October 2022 (worst case), then the delivery of cassettes would be delayed by 4.5 months unless cassette assembly is accelerated.

DC-DC convertors -- The ASICs for the DCDC system are scheduled to become available 18 months before they are needed for the CE, so we expect no impact on the schedule even if significantly delayed.

Table of Deliverables

WBS	Item	Threshold Quantity inc Spares	Objective Additional Quantity
402.4.3	Si Full size sensor wafers	7854	
402.4.3	Si Odd-size (Full-size Equivalent/wafers)	1265	58
402.4.4.3 & 402.4.4.4	Si Standard Modules	9753	387
402.4.4.5	Si Odd-size CE-E Modules	3354	135
402.4.4.5	Si Odd-size CE-H Modules	2450	102
402.4.5.1.2	Si Motherboard engine	2970	115
402.4.5.1.2	Si Motherboard wagon	5940	230
402.4.5.1.3	CE-H Cassette Interface	616	22
402.4.5.3	CE-H Cassettes	528	22
402.4.6.2	Wrapped and QC'd scintillator tiles	140k	10k
402.4.6.3	SiPM photodetectors	132k	10k
402.4.6.4	Assembled tile modules	1248	156
402.4.6.5	Scintillator motherboard assemblies	1008	42
402.4.7.1	ECON-T	41k	
402.4.7.1	ECON-D	41k	
402.4.7.2	LV Power Supply		1200
402.4.7.2	HV Power Supply		2150

MIP Timing Detector

1. BTL and ETL Assembly: How do you qualify the assembly procedures? Also, please describe the QC testing and cold testing of trays and modules.

Rephrased: Four US assembly sites are used in the MTD, what are the plans for site qualifications, tools development and testing across the sites (nationally and internationally)?

BTL assembly:
All components for the BTL tray assembly (crystals, SiPM, FE cards, cooling trays) are delivered tested at their respective production/reception sites.
At the assembly site the first step is the mating of the SiPMs with the LYSO arrays. After this step the sensor module is tested. Testing can be done easily since the **natural radioactivity of LYSO** creates sufficient signal to measure the S/N of components. The following steps are mounting components on the cooling tray. **After each step the same test procedure will be carried out.** At any step, if a component is found to be faulty it can easily be replaced. **After a full tray is assembled it will be cooled down in a cold box and the full tray is tested in the cold.**

Qualification of the assembly center:
Each assembly center **will build a prototype tray**. This will serve to qualify the assembly center procedure. All assembly centers are deeply involved in the tray assembly R&D and exchange technical details on procedures, tools and test protocols. All assembly center will **follow a uniformized, agreed upon procedure**. Travel cost for visits between assembly centers are included in the BoE to allow hands-on training and synchronization of procedures.

ETL assembly:
For ETL UNL is the pilot site, but all assembly sites will be qualified for production through a sequence of prototyping and pre-production activities. Each site will **demonstrate the assembly procedure in a few steps of increasing throughput rate**: 1, 10, 50, 100 modules/weeks. The last step represents the full throughput capacity that we will achieve at U.S. assembly sites. These demonstrations will be carried out with mechanical components with the proper geometry during 2020 and 2021.

Additionally, both U.S. sites will **demonstrate assembly of functional modules prior to production**. Initially, each site will produce a few modules with the ETROC2 prototype in 2022. Each site will then participate in a pre-production run in 2022-23 during which it will produce 200 two-sensor and 50 one-sensor modules.

Sites are required to successfully **complete all of these activities to demonstrate assembly quality and rate** prior to being qualified for module production which beings in late 2023.

QC testing of production modules is comprised of several activities. Upon reception, all components will be inspected and the flex tested. During assembly, all modules are **tested to verify electrical connectivity three times**: 1) after wire bonding the ETROC-LGAD (ELA) assemblies to the flex, 2) after encapsulating the wire bonds, and 3) after the module is fully assembled. Additionally, after assembly, all modules will go through 1) a full electrical test stand where all required functionality will be demonstrated and 2) testing of basic functionality during thermal cycling.

A subset of each batch of production modules **will also undergo burn-in and evaluation in a test beam before and after irradiation**. Some modules will also be reserved for testing to failure. The same set of activities described for production modules will also be performed for pre-production modules.

> *2. **Uncosted labor:** Can you tell us more about the types, profile, locations, and responsibilities associated with the uncosted labor in the project? Will these needs appear in the SOWs?*
>
> *Rephrased: There is a large fraction of uncosted labor in the MTD part of the project. How is the uncosted labor committed, tracked and utilised? do you have a breakdown of contributed labor (student, PostDc, faculty, scientists) per activity?*

Chris has prepared summary slides with this information:slides

Regarding the SOWs, the standard approach in preparation of SOWs is to include a list of the uncosted labor dedicated to this project, including faculty, postdocs, and graduate students, which should be listed with a cost of $0 to USCMS.

> *3. **ASIC engineering:** How will you ensure that your ASIC engineering needs are satisfied given the situation in the FNAL ASIC group?*
>
> *Rephrased: ETROC ASIC for the ETL part of the project is on the critical path. Do you have sufficient resources for the ASIC design for the full duration of the project?*

There are currently three main ASIC projects at Fermilab: Cold electronics for **DUNE**, CMS **ECON**, and CMS **ETROC**. The next round of DUNE cold electronics (final version) chip submission is scheduled to be done by early 2020, while the ECON submissions are scheduled to be done by summer 2020. The plan of the ETROC2 development has taken this into consideration. The ASIC engineering resource sharing has been carefully monitored by Fermilab ASIC PMG, which has helped to resolve potential conflicts. In fact, Fermilab now has three ASIC designer positions opened to further alleviate potential future resource conflicts.

> *4. **Testbeams/irradiation:** What are your testbeam and irradiation needs in the coming years, through the era of destructive testing of production-era devices?*
>
> *Rephrased: Test-beams and radiation testing are very important for testing ASICs and*

assemblies of ASICs+sensors. What are the plans for beam tests and irradiation tests for BTL and ETL? Is the FNAL test-beam facility sufficiently supported for the duration of the project until at least the end of pre-production?

For SiPM irradiation, we will need one round shortly after the preproduction SiPMs arrive at CERN (July 2020), then another shortly after the first production SiPMs arrive at CERN (July 2021). For completeness, another irradiation towards the end of the production in order to compare radiation tolerance of SiPMs from throughout the entire production cycle. We have had, and will continue to have, access to the JSI facility in Ljubljana.

Testbeams BTL:
We plan 2 to 4 test beams per year at FNAL, using the dedicated timing infrastructure which we have been using since 2017. Further details of this setup are given below.
The current plan for BTL test beam campaigns are :

- TB Feb 2020 and May 2020 : module tests with TOFHIR1, Mechanics prototype close to final, first TB to merge FE, sensor, and mechanics R&D
- TB Nov 2020: TOFHIR2v1 with irradiated sensors with cooling of SiPMs
- TB 2021: TOFHIR2v2 + full RU, all irradiated

We plan to test sensors and modules 2-3 times a year, and we have an FNAL experiment dedicated to CMS timing detector (T1409), which has received 3-to-4 dedicated periods every year for the last 3 years, each 2-week long. The ETL has a permanent setup at the test beam at FNAL (FTBF), where modules are being tested and evaluated, complete with a DAQ and cooling. The setup has been used extensively over the last few years. Since this is a permanent setup in FTBF, measurements can also be done in a parasitic mode concurrent with other experiments, in case of component delays.

For irradiation of modules and prototypes we have been relying on the Ljubljana nuclear reactor facility, which we have been performing through our collaborators from Torino. We are now in the process of transitioning to using the Rhode Island facility for neutron irradiation, and we plan to perform 2-3 rounds of irradiation per year. Additionally, FNAL is in the process of commissioning the proton irradiation facility, which we will use for proton irradiation of modules, and for SEU studies of the ETROC.

Below is the list of currently planned test-beam campaigns, each of which is preceded by neutron or proton irradiation of modules.
TB Feb 2020 and May 2020: tests of ETROC0+LGAD
TB Aug 2020 and Oct 2020: tests with ETROC1+LGAD
TB Mar 2021: test beam ETROC1-modules
TB Jan 2022 and Feb 2022: test beam with ETROC2 based modules
TB Nov 2022 and Dec 2022: test-beam with pre-production modules

5. Can you clarify what is the targeted time resolution, i.e. 30-60 ps?

We say **30-40ps at beginning of life**. This is because although we have demonstrated the inherent resolution of LYSO bars+SiPMs and LGADs provide 30ps or better, these studies were 1. single-sensor, 2. using a different ASIC than in the MTD, and 3. using a different reference clock than in the final system. Hence we say 30-40ps to cover the contributions to the resolution from these other sources.

We also say **<60ps through 3000/fb**, since we have shown from simulation that MTD will achieve this for 3000/fb assuming a nominal fluence with a 1.5 safety margin and the simulation results are consistent with the irradiated results we have on individual devices.

In the **KPPs we just say <60ps for simplicity**. This follows guidance from a comment we received at the Director's Review. The benefits of resolution of 60ps are still clear -- there is no performance cliff that we encounter even up to a time resolution of 70-90ps.

Trigger/DAQ*

**These questions were sent by email on 10/21 to be answered by during the breakout*

1. The TDAQ schedule is remarkable in that there are large numbers (up to 15) of tasks with identical start dates & lengths - including critical path activities. There are also multiple groups of tasks with identical task durations. Is the work planned the way it will be performed? To what extent are these durations placeholders? How have the task lengths been defined? How are they likely to evolve and has the possible effect on the schedule been considered?

This was discussed in the breakout session with the drilldown team. It was explained how groups of tasks were linked to satisfy simultaneous hardware, software and firmware milestones, with sets of software and firmware features planned to coincide with different stages of hardware development. Suggestions were made for how to make the software and firmware tasks more explicit about what is included in each milestone, and how to track their value.

2. Slide 17 in your plenary talk quotes a 21% uncertainty. This seems low given that the project is only ~60-70% complete in terms of final design and ~45% complete in terms of detailed design according to the Design Completion Definition (13417). The pie chart on slide 17 shows 84% L4/L5 (25-40% uncertainty) which also suggests a higher uncertainty. Please go into these assessments -- and the risk evaluation -- in detail during the TDAQ breakout session.

This was discussed in the breakout session. 21% corresponds to the ratio of EU/BAC. When the uncertainty is more appropriately calculated as estimate uncertainty and risk contingency together as a fraction of the work to go (EU+Risk)/CTG= (2.4M+1.11M)/7.8M = 45%, it is a more comfortable number. The slides will be updated to make this more clear.