

#	Risk Event Description	Potential Ramifications	Impact of Risk	Likelihood of Risk	Risk Level	Strategy
Management Risks						
1	Schedule delay for PPD&TD rollout phase, or some other phase, due to one or more of the other risks/events identified. Possible delay of milestone date for entire lab on FTL for both time and effort reporting.	<ul style="list-style-type: none"> Impact on Finance/Accounting if parallel effort on electronic system and paper system has to be maintained. DOE negative attention on missing a milestone that is used in a PEMP metric. FRA board attention and displeasure on missing a milestone. Embarrassment at slipping the schedule yet again. Delay in lab getting a full accounting of where effort is being applied. 	Not clear – could be anywhere between low and very high.			<ul style="list-style-type: none"> Develop a clearer understanding of the actual impacts of slipping the schedule of one or more phases of FTL rollout. Examine possible changes to the schedule that nonetheless result in the entire lab being cut over to electronic timecard and effort reporting around the same end-date as in the original schedule and thus minimize negative impact on performance measure scores. Re-examine strategy for going live with both timecards and effort reporting at the same time. (needed to address an audit issue) Work with Finance/Accounting to determine what would minimize their workload and take into account any other milestones or deadlines that affect them. If one of the factors in a schedule slippage is the changing landscape due to the American Recovery and Re-investment Act then examine whether explaining this might alleviate some of the potential impacts and ramifications.
2	The Computing Division Head/CIO (acting) now carries ultimate responsibility for the FTL service being ready for deployment and for it being acceptable and usable by Finance/Accounting, end users, division administrators, supervisors, managers, and project managers.	<ul style="list-style-type: none"> CD Head/CIO(acting) has had little time to assess all aspects of the FTL project and the FTL service and to become comfortable with signing off on the FTL service as ready for expanded use over that currently provided to the Sections and the Directorate. She is unlikely to be willing to sign off on accepting all of the risks identified and their mitigations without additional scrutiny. 	<p>Low – if the Management Risk of modification of schedule can be mitigated and the results of a review lead to reduced risk later</p> <p>High – if the Management Risk of schedule modification cannot be mitigated.</p>	High		<ul style="list-style-type: none"> Some time is needed to sign off on the technology and process deliverables of the FTL project and to restructure the operational work under an FTL service owner. The operational readiness of the FTL service and the state of customer readiness to adopt and use the service needs to be assessed by people other than the FTL project team themselves. A schedule delay of at least one month is needed just to do this and to develop corrective actions for those risks which the CD head/CIO is not yet comfortable accepting. Apart from the risks identified elsewhere in this risk register there are additional potential risks that need to be examined such as <ul style="list-style-type: none"> Limited browser support Need for terminal server client code and impact on desktop support services Interplay between computer security requirements to patch systems and browsers, but Oracle EBS system expectations to control browser versions in a homogenous environment. Password reset being handled by a one-person help desk during business hours with no plan yet developed (nor load understood) to extend helpdesk support during initial stages of adding new users Confusing and duplicative business processes for maintaining employee information and authorization information. No experience with managers and project managers using the FTL effort data for management purposes, replacing existing mature effort reporting and management business processes. Monitoring and metrics for availability and performance of the FTL service and yet-to-be clarified agreements on service restoration outside normal working hours and response to users outside normal working hours (who will likely call CD helpdesk if they are having trouble entering their timecard). Requirement (or lack of) for FTL service as part of disaster business process continuity scenario – clear understanding of what can go to paper processes in a disaster scenario and what cannot.
Business/Operational Risks						
1	Adding additional users to the system could adversely effect the time required to run the payroll BEE (Batch Element Entry) processes	<ul style="list-style-type: none"> These processes are run weekly as part of current payroll procedures and one of the process (validate and transfer to BEE) run-times was increasing exponentially, impacting an already tight timeframe for weekly payroll processing. As part of the ATG upgrade, we identified and applied an Oracle Patch that resolved this issue (confirmed in our parallel testing) 	High	Low	Low	<ul style="list-style-type: none"> Applied and tested Oracle patch to address the issue
2	Operational risks in divisions related to effort reporting review	<ul style="list-style-type: none"> Currently, in the divisions, supervisors typically approve leave sheets and project managers are given an opportunity to review project hours before they are sent to Project Accounting. In FTL, we only have the ability for one approver (the person who best knows what hours were worked/not worked) and don't have a formal systems review process for the project managers prior to posting these transactions to Project Accounting. In FTL, timecards need to be entered by 10:00 am on Monday and approved between 10:00 am and 11:00 am (although, even with the sections we see approvals late into the afternoon on Monday and expect the same for the divisions). Once the timecards are all approved, two processes are run that "move" the timecard data: <ol style="list-style-type: none"> from the Timestore to Payroll (BEE processes) 	Medium	High	Medium	<ul style="list-style-type: none"> Discoverer queries have been created that will give the project managers the ability to review hours charged to their project on a weekly or monthly basis Existing review processes that use Project Accounting or COBRA will still be available to the project managers Development of alternate division-specific and project-specific business processes for review of effort and corrective actions after the fact, rather than prior to posting to PA as currently implemented Development of division-mandated timeframes for actions and clear expectations for roles of timekeepers and timecard approvers.

#	Risk Event Description	Potential Ramifications	Impact of Risk	Likelihood of Risk	Risk Level	Strategy
		<p>2. from the Timestore to Project Accounting.</p> <ul style="list-style-type: none"> Currently, we are running the BEE processes for monthly employees on Wednesday every week, and running the PA interface Monthly on the Tuesday after the third Sunday (this interface will be run weekly after the entire lab is on FTL.). The risk that these tight timeframes creates is that the project managers will not be given adequate time to review the project charges before they are sent to PA. This risk is even greater once we implement nonexempt employees because the BEE processes for nonexempt employees will need to be run on Tuesdays to meet the deadline of end of day Wednesday to get a payroll file to our bank for processing 				
3	Lack of staff redundancy in Finance/MIS/field	<ul style="list-style-type: none"> The resource levels across Finance, MIS, and the field that need to be involved in FTL processes lack redundancy, and are at risk if "key " individuals leave the lab 	High	Medium	Medium	
4	Lack of funding to sustain a functional consulting resource (Gene)	<ul style="list-style-type: none"> One of our lessons learned from the initial implementation was that we needed to engage a full time functional consulting resource for the duration of the project. Currently, the level of funding allocated to the project provided a budget for a full time resource from October thru January, and then two days per week from February thru August. We have two risks related to this issue.: <ol style="list-style-type: none"> The consulting company may find a full time engagement for this resource and will pull him from our engagement. This resource is critical to the testing and implementation of the non exempt timecard, and two days a week is not enough time for him to adequately support the testing and implementation activities 	High	High	High	<ul style="list-style-type: none"> Work to modify budget across MIS and other CD organization units (in conjunction with the budget office) in order to maintain adequate functional analyst resources. Begin cross-training in FTL/Accounting functional area.
5	Risk of requirements change from external or internal entities	<ul style="list-style-type: none"> We are always at risk that business requirements may change and impact the availability of the resources that are allocated to existing project activities. We also have the risk that an external requirement is mandated that impacts the project. We can mitigate the risk for internal changes (centralized authentication, new reporting requirements, etc.) by enforcing a strict change control process, but have little control over the changes to requirements (EVMS cert) that are mandated from external sources (OECM) 	High	Medium	Medium	
6	Loss of credibility if we start moving groups off of FTL	<ul style="list-style-type: none"> Risk of losing credibility with the divisions that we will be rolling out if the organization implementing the change is no longer using FTL to enter their own time 	High	High	High	<ul style="list-style-type: none"> MIS organization will continue to use FTL – this risk does not exist – unless of course the FTL system cannot deal with particular individuals from CD being on FTL. Extracting the MIS effort data from FTL and using it together with existing data may actually help mitigate some of the other business/operational risks.
7	Risk that training for TD and PPD will need to be re-executed	<ul style="list-style-type: none"> If the project dates are pushed the just in time training currently being executed for TD and PPD will be lost/forgotten, and we will need to retrain them at a latter date 	High	High	High	<ul style="list-style-type: none"> Only the timecard approvers and those doing maintenance of employee data need just-in-time training. They may need more than one session of training, or individual attention and support anyway.
8	Loss of credibility if we move the dates for PPD and TD	<ul style="list-style-type: none"> If the project dates are pushed, it could result in the loss of the project's credibility and TD and PPD upper management support, which has already incurred a significant cost (resources) allocated to making the March 30th date 	High	High	High	<ul style="list-style-type: none"> PPD management does not see a schedule slip as significant. Actually increases credibility if their state of readiness improves. TD management attitude is not known at this point.
9	The American Recovery and Re-investment Act of 2009 is very likely to result in significant funds to Fermilab that must be committed rapidly and also tracked and reported on in ways that have never been done before.	<ul style="list-style-type: none"> It is likely that the information management required to meet reporting requirements will result in calls on MIS (and other CD) staff to rapidly help support new business processes that collect and store new data that must be combined with core data from the procurement and project accounting system to satisfy federal reporting requirements. This is a specific and likely risk related to the 	Not yet known But potentially high or medium	High		<p>Use the next month to</p> <ul style="list-style-type: none"> Mitigate management and technical risks wherever possible Develop strategies for dealing with schedule change including potential reprioritization of other projects and initiatives.

#	Risk Event Description	Potential Ramifications	Impact of Risk	Likelihood of Risk	Risk Level	Strategy
		Business/Operation risk (5) identified of "Risk of requirements change from external or internal entities"				
10	Changes for CERN/FRA Calendars	Issues related to costing the CERN and FRA people were identified late in the development cycle, and may not be ready for the 30/30 cutover.				<ul style="list-style-type: none"> • Dedicate testing resources to test the changes related to CERN instead of testing the nonexempt timecard process (potentially causing a delay in the end of the nonexempt system test) • Go live without CERN employees and have them start later in the month

Technical Risks						
0	Computer Security governance and policy have imposed a requirement that the FTL service use central authentication prior to full deployment to the entire laboratory. This is an issue outside of the FTL project/service itself but clearly impacts the work to be done and therefore potentially the schedule for rollout of FTL to additional groups of people in the lab.	<ul style="list-style-type: none"> Meeting the requirement will necessarily create a delay between some phase of the rollout schedule. Working on a technical solution now may take so much time away from mitigating some of the other risks that rollout of PPD&TD on schedule would anyway not be possible. 	High	High	High	<ul style="list-style-type: none"> The sooner the upfront work is done on understanding the length of the delay needed to rework parts of the technical underpinnings, the sooner a revised schedule can be made that produces the same, or close to the same, end-date for a complete rollout of FTL to the whole lab. Delaying deployment to further parts of the lab for at least one month will allow time for staff to work on designing and implementing the technology that meets the authentication requirement (as well as mitigating other risks) without having to simultaneously support and improve a service that PPD and TD begin relying on. <ol style="list-style-type: none"> Deploying to PPD and TD in April, using existing application level authentication+ authorization carries the risk that MIS development staff will spend a great deal of time responding to issues generated by new modes of usage (by divisions and projects) that have not been exercised up to now with division-level business processes that have not been fully developed. Risk is then of longer schedule delay than the minimum needed to deploy a central authentication solution. Deploying to only TD in April, using existing application level authentication+authorization carries the same risks as (a) but may be slightly mitigated by TD having a better understanding of their own business processes. Risk is still that of longer schedule delay than the minimum needed to deploy a central authentication solution.
1	Errors by inadequately trained administrators leads to system downtime	<ul style="list-style-type: none"> Due to staff losses over the last year, budget limitations affecting funding professional services, and long recruiting intervals for new staff, the MIS department DBA team has been understaffed and unable to build necessary skill and support redundancy to adequately protect against the risk presented by the absence or loss of a single member of the DBA team. The absence of key individuals increases risk to the system, since if complex system administration activities need to be executed during the absence, sufficient skills are not available to perform the activities without introducing a higher risk of error. The ever increasing demands, hours, and plummeting morale of the MIS department, combined with ongoing budget and programmatic uncertainty for the laboratory, increases the likelihood of permanent loss of key staff. 	High	Medium	Medium	<ul style="list-style-type: none"> Potential loss of staff for whatever reason can only be addressed by <ol style="list-style-type: none"> Cross training additional staff over a period of several months Making sure risks are understood and mitigated where possible Making sure schedule is reasonable
2	System change results in processing errors or downtime	<ul style="list-style-type: none"> A number of incidents in the past year or so have shown that significant, critical application problems have escaped detection during testing cycles, only to be discovered once the associated changes have been moved to the production environment. Some of these problems have been severity 1 cases where business operations have been threatened. Only the combination of extraordinary staff efforts and luck have these issues been addressed quickly to prevent serious impact to business operations 	High	Medium	Medium	<ul style="list-style-type: none"> An operational readiness walkthrough/review of various processes and procedures, by "outsiders" not familiar with the system could possibly point out some hidden issues
3	Insufficient investment in maintenance allows minor problems to grow into major ones which threaten application availability or integrity	<ul style="list-style-type: none"> The reduction in resources in the MIS department coupled with high priority project implementations and commitments that must be met results in reduced maintenance efforts, reduced staff training, and trending towards break-fix quick solutions, rather than root cause analysis. This increases the risk that insufficient time and attentions are dedicated to seemingly minor problems and that these problems can then grow to threaten the system overall or that leading indicators of a major problem are overlooked. Recent examples of this include login problems in EBS due to middle tier communications problems and database process exhaustion issues after 10g RDBMS upgrades 	Medium	Low	Low	<ul style="list-style-type: none"> An operational readiness walkthrough/review of various processes and procedures could possibly help Cross-training additional staff over a period of several months may reduce pressures and increase depth of expertise available to deal with difficult problems.
4	Database server memory exhaustion	<ul style="list-style-type: none"> Initial performance data seemed to indicate the database server might be encountering a memory problem with our current, normal operational load - this was indicated by the fact the DB server is doing memory swapping. However, on further analysis, it appears the DB server has about 8GB of headroom (of its 12GB total memory) and that, while it is swapping, the swapping isn't extensive and may be due to a combination of little, consistent process demand for memory and the extreme load created by the hanging logins problem two weeks 	High	Low (if short term steps are taken)	Low (if short term steps are taken)	Remediation steps are: <ul style="list-style-type: none"> SHORT TERM <ol style="list-style-type: none"> Reboot bssopsf01 cleanly to start observations on swap usage from a clean boot. Acquire additional memory modules to allow us to expand physical memory from the current 12GB to 16GB and to have sufficient memory on hand to expand to the maximum memory capacity of 32GB. This will be accompanied by purchasing sufficient memory for bssopsf02 (QA) to expand to 16GB. Est. cost is \$1600 for this memory. Reconfigure bssopsf01 & 02 to use the hugemem OS kernel, which will permit the system to use beyond 16GB of physical memory (this change is transparent to the applications and just puts us in the situation that we can use >16GB RAM, but we will run this hugemem kernel with <=16GB RAM initially) Continue investigation of Linux tuning recommendations for RDBMS 10g on RHEL 4 to see if additional tunables should be adjusted on

		back (that appear to have pushed the server deeper into swap and it has remained there since). But, given this data, I have to consider this memory issue as a problem which may worsen as we add users.				<p>the OS</p> <p>5. Adjust the DB node to use gigabit networking (it currently is configured for 100Mb) - the hardware is in place - this is just an OS setting change.</p> <ul style="list-style-type: none"> LONG TERM <ol style="list-style-type: none"> Spec and acquire a "beefy" third web node for PRD that will enable us to choose to move the current web & forms node from bssopsf01 to a separate box, thereby freeing about 2-3GB RAM and associated processing from the DB node. Cost ~\$18K for this server (DL580G5 with dual quad core processors and 16GB RAM) Evaluate the existing EBS server network switch for possible replacement to insure it can handle maximum required traffic between RDBMS and web nodes- the current switch is a low-end switch which is likely to be saturated if the gigabit networking links between the web nodes and the RDBMS node were to be heavily utilized (not currently happening). Guessing a replacement switch would run about \$25K, but this is rough guess and depends on what network configuration is adopted (single switch in current topology or an additional, second switch in a revised topology). Upgrade the FCC SAN switch to 8Gb capable switch and replace the SAN adapters in the DB server with high speed adapters to minimize the chance of an I/O bottleneck. This is basically applying the planned SAN upgrades in FCC that we have done in Wilson Hall, plus doing the server side upgrade to maximize advantage of the new SAN fabric speed for these servers. The SAN switches are \$49K. The server adapters are about an additional \$7K (covering DEV, PRD, and QA) Plan for a future upgrade to EBS12, which is a prerequisite to moving our current EBS single node DB/forms/concurrent manager forward to a x86-64 Red Hat 5 platform using its current configuration Investigate the effort involved in separating the concurrent manager and RDBMS nodes (an EBS split configuration) to allow the DB node to be fully dedicated (and tuned and optimized) to the RDBMS. Executing this separation would likely incur significant modification to application customizations (which, I believe, assume the concurrent manager node and RDBMS node are one and the same) and greatly complexify our database refresh process. The new node acquired above would be the intended target for the concurrent manager, so this is why its hardware would be more capable than the current web nodes. Plan for a hardware replacement of the DB server node with x86-64 architecture running 64-bit RHEL 5, assuming an EBS 12 upgrade is done before this hardware upgrade (otherwise, we can't migrate the existing EBS DB node architecture to 64-bit RHEL 5). This hardware would be capable of running up to 64GB RAM.
5	Middle tier excessive load	<ul style="list-style-type: none"> While there is no sign of a performance issue on the middle tier for any current operations, with the exception of when we were having hung login problems, the nature of the architecture dictates this likely to be a pinch point as user count increases. Even under the heavy load of the hung login problems, the middle tier nodes showed very little load. It is expected each middle tier should be able to handle from 200 to as many as 600 concurrent users. These nodes are processing nodes, so load will increase on CPU and memory, but not much on I/O. These nodes have maximum CPU installed (2 sockets with dual core CPUs - so, effectively, 4 CPUs) and have 12GB of 16GB max memory installed. 	High	Low (if short term steps are taken)	Low (if short term steps are taken)	<ul style="list-style-type: none"> SHORT TERM <ol style="list-style-type: none"> Work with apps team to get some realistic monitoring of concurrent user load to understand current user load as a basis for extrapolating resource load for future user load. Investigate and tune Java/JDBC settings on middle tiers based on Oracle Support documentation and whitepapers. Investigate effort/feasibility of being able to quickly reallocated existing QA middle tier nodes to PRD if needed LONG TERM <ol style="list-style-type: none"> If increased load appears to be generating rapid increase in memory usage, purchase additional memory to grow nodes to full 16GB RAM (est. cost <\$5K) Spec and acquire a "beefy" third web node for PRD that will allow us to add a third web node. This is the same box that is describe above (under LONG TERM steps for DB node) at an estimated \$18K. Evaluate the existing EBS server network switch for possible replacement to insure it can handle maximum required traffic between RDBMS and web nodes- the current switch is a low-end switch which is likely to be saturated if the gigabit networking links between the web nodes and the RDBMS node were to be heavily utilized (not currently happening).
6	MIS network firewall load	<ul style="list-style-type: none"> The MIS network firewall is a single chokepoint for all EBS network traffic. And, in fact, a round trip transaction from an EBS client through the middle tier to the DB node traverses the firewall 6 times, so dramatic increases in input network traffic to EBS are magnified here. However, the network firewall shows no signs of any load problem. Most likely concern would be memory, since this box is minimally configured for memory. At most, a tweaking of firewall configuration settings to support more concurrent network sessions might be needed. Worst case, we would have to consider a significant hardware upgrade to provide increased processing power for the firewall or, probably a better approach, would be to either remove the firewall from between the middle tier nodes and DB node or place a separate, dedicated firewall node in this location. 	High	Low (if short term steps are taken)	Low (if short term steps are taken)	<ul style="list-style-type: none"> SHORT TERM <ol style="list-style-type: none"> Purchase additional memory for firewall and backup to bring up to 8GB (from current 4GB) - cost around \$1K.
7	MIS EBS load balancer	<ul style="list-style-type: none"> The MIS load balancer is another single chokepoint for all EBS network traffic destined for the middle tier nodes (time.fnal.gov). Again, there is no sign of a load problem here. Some memory might be added to this minimally configured server, but that is unnecessary unless we see a significant increase in memory usage with future load. 	High	Low	Low	<ul style="list-style-type: none"> If the processing load becomes too much for this system, the most likely solution would be investing in a commercial, dedicated load balance appliance (CISCO, etc.). I am guessing such a load-balancer migration would have to be coordinated with the CD networking group, so I can give no reliable estimate of cost
8	SAN storage array	<ul style="list-style-type: none"> The database server I/O depends on the SAN storage array I/O. With the currently in progress EVA upgrades, 	High	Low	Low	<ul style="list-style-type: none"> If the SAN storage array performance appears to be problematic, we would need to consider a large, rebalancing of our DB server to storage architecture.

		<p>we are likely to move the maximum I/O bandwidth capacity of the SAN arrays well beyond what our current database server architecture can generate.</p> <ul style="list-style-type: none"> As mentioned above, we may need to upgrade the SAN fabric (server to array interconnectivity) to insure I/O isn't restricted by this. 				
9	Residual risks identified in Business Systems MA security plan CSP- MA-991	<ul style="list-style-type: none"> 18 residual risks identified (for all business systems) and mitigated to at least the "Low" level. 			Low or better	<ul style="list-style-type: none"> Make sure that the FTL service and its usage by the entire lab does not change any of the controls needed, or any of the assumptions about residual risks, other than the authentication issue already noted in (0).