Crosstalk Article [Targeted Issue]: Strategizing with Agility, May 2025 Issue

Risky Business: A Comprehensive, Agile Approach to Risk Management

By Stephen Wilson and Stephanie Grzenia

The Problem with Risk Management

While many Department of Defense (DoD) programs today have an active risk management (RM) framework in place for their cyber-physical weapon systems, most of them are tuned for long-term risks that require significant time and effort to manage. Even if programs are doing a lightweight Agile risk management approach, such as ROAM, which generally has lower costs, it is often siloed from the other traditional RM system [1]. To achieve a comprehensive RM process, organizations need a solution that addresses each type of risk — long- and short-term — because a realized risk can have profound implications for an organization, its clients, and beyond. For example, Equifax, in 2017, provided a cautionary tale when it didn't update a key security patch — a risk that was known but not properly addressed, and hackers accessed personal data of an estimated 143 million Americans [2].

To protect against this sort of breach and ensure effective RM, we recommend that the traditional approach to managing long-term risks be combined with such Agile RM practices as Risk ROAMing, which organizations like Scaled Agile, Inc. have used to address short-term risks. The DoD references the possibility of combining approaches in its seminal RM guidebook the Risk, Issue, and Opportunity Management Guide for Defense Acquisition Programs (the RIO Guidebook) [3].

Aligning with the RIO Guidebook, our recommendation is that a more efficient RM system can be made by using a lightweight approach like Risk ROAMing as the first line of defense against short-term risks and as an iterative and incremental approach to address long-term risks. Additionally, organizations can benefit from the data and insights produced through the Agile process to inform — and adjust — the approach to managing long-term risks. We recommend using Agile best practices as an empirical approach that favors rapid incremental reduction of risk. This approach allows programs to more completely close the smaller risk mitigation steps on a frequent Agile cadence. These risk mitigations steps are decomposed from the larger risk statements. Using this concept, the program can retire technical debt and risk mitigation steps in a timely manner.

Figure1. Traditional RM and Agile RM approaches, like ROAM, both have benefits and challenges.

In this article, we define both types of RM approaches and explore the benefits and challenges accompanying each. We then explore how integrating them creates a robust RM system that effectively mitigates both short- and long-term risks.

What Is Traditional RM?

Most readers are likely familiar with traditional RM [4]. Different companies and teams have their own ways of documenting and tracking risks; however, the typical process entails documenting crowd-sourced risks in an RM tool where each risk is given a unique ID and a conditional "if/then" statement. The risk manager then assigns a criticality score, commonly a two-dimensional value based on the probability of occurrence and the severity if the risk is realized. If available, a mitigation strategy is then assigned to the risk, and any additional background information is documented in the RM tool. The risk is then given a priority ranking and assigned to a risk manager who works with stakeholders (referred to as a risk review board) on a scheduled basis to re-evaluate the effectiveness of the risk's ongoing mitigation strategy and probabilities of occurrence and severity. The frequency with which the risk(s) should be reviewed, who needs to review them, and the scoring rubric have nuances that make these elements unique to each traditional RM program.

Traditional RM is not without merits though. Tools and process guidance abound for this type of RM because it is necessary to identify, capture, manage, and track the long-term risks that program offices discover at all phases of the product life cycle. When multiple risks exist and are tracked, the administrative cost of scoring the risks, prioritizing them, updating them in an RM database, and periodically reviewing them can prevent the program from being caught off guard.

Challenges with Traditional RM

Many traditional RM frameworks often follow a waterfall approach, requiring extensive up-front effort to define and design a mitigation strategy for the entire risk life cycle. However, because the pace of work doesn't slow, many RM teams never catch up on evolving demands, and unexpected changes can provoke a need to rework the initial plan. Due to these challenges, it often takes many months or years to see a risk all the way through mitigation to resolution or realization; therefore, significant administrative costs arise. Alongside these administrative costs come mental and emotional costs as these risks require maintaining a big-picture perspective. These costs can result in a reluctance to document and address long-term risks. Crucially, because risk mitigation is often disconnected from the teams who are doing the work, there may not be enough opportunity for consistent or timely feedback generated from the solutions that are implemented in an attempt to mitigate or manage a risk. Applying Agile approaches like ROAM for short-term risks and decomposing long-term risks into smaller pieces of work for iterative and incremental burndown addresses or reduces many of these challenges.

What Is ROAM?

Risk ROAMing is an RM philosophy built for Agile development and promoted by frameworks like the Scaled Agile Framework (SAFe). Risk ROAMing focuses on identifying and managing risks that could prevent a team from achieving their commitments for the current cadence-release time box. It is particularly useful for organizations that engage in big room Agile cadence release planning, which is when multiple Agile teams plan their work for the upcoming planning cycle, a period that is usually about three months long. (SAFe calls this PI Planning.)

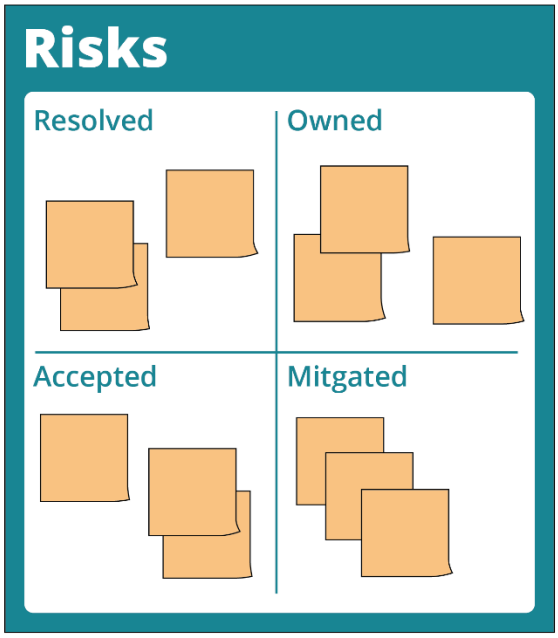

Figure 2. Risks can switch between and move through the four ROAM categories throughout the RM process.

In big room planning events, teams often rapidly identify risks over the course of a couple days as they plan their work for the planning cycle. On the final day of the planning event, the risks are "ROAM'd" into one of four categories. These categories are states of being through which the short-term risks can switch and move throughout their life cycles. These four ROAM categories are:

- Resolved — A risk is resolved when the Agile teams agree that the risk has been addressed. Risks identified early in big room planning events may become resolved later in planning when more information is available.
- Owned — A risk is owned when a mitigation strategy is believed to exist, but there is not enough information available at the time of the ROAMing event to identify it fully. Therefore, a designated individual or team takes ownership of the risk and is responsible for doing more research and potentially developing a mitigation strategy later. The risk owner's goal is to move the risk to a mitigated, accepted, or resolved state as soon as possible.
- Accepted — A risk is accepted when no further action will be taken to resolve the risk because mitigation is either too costly or not available. (Accepted risks are often referred to as "watch items" in traditional RM.)
- Mitigated — A risk is mitigated when Agile teams have identified a mitigation strategy for the potential event to reduce the likelihood or consequence and have incorporated that strategy into their committed plan. If the mitigation strategy is carried out as planned, then the risk can be thought of as resolved or accepted.

Ideally, the risk is eventually resolved or at least effectively mitigated to the point that it becomes accepted. During big room planning, the teams collectively agree on the risks and their associated categories.

One of the main advantages of ROAM is its efficiency; it can be managed at a team or individual level, making it less administratively burdensome for the organization. Additionally, Risk ROAMing improves team-level understanding of risks and emphasizes their resolution by allowing team members to identify and discuss them during big room planning, ensuring visibility on the program boards. This democratization of the risks increases visibility, fosters greater buy-in, and focuses the Agile teams on mitigation and resolution.

Risk ROAMing also reduces the time spent discussing risks compared to traditional processes because the scope of the discussion focuses on a relatively short time horizon. In well-run Risk ROAMing events, teams can identify and process 5 to 10 risks with a room full of people in less than 30 minutes, and very often even less than 15 minutes.

This efficiency is possible because the risks are generally smaller in scope, confined to the planning cycle, and their potential impacts are almost always the same: if realized, a given risk will prevent the team from accomplishing their planned work during the time box. Since these risks have a much shorter life cycle, they do not require the additional analysis and discussion to quantify their probability and severity. Furthermore, ROAM'd risks generally are not plotted on a risk assessment matrix nor prioritized because their mitigation strategy is either incorporated into the plan or not.

Challenges with ROAM

ROAM is often criticized for the confusing nature of its category naming and mapping. Different publications have slightly different nuances for all the terms. This discrepancy can make categorization and team collaboration difficult. For instance, it is easy to understand why someone might think that "mitigated" means that the risk has been fully resolved, rather than meaning that the affected members/team have planned for a mitigation strategy. Solutions to avoid this misunderstanding are to ensure the definition of each category is discussed before ROAM events and that the definitions are displayed and easily visible to the participants during ROAMing.

Another challenge with Risk ROAMing is the difficulty that often arises when distinguishing between a risk and a dependency. Like risks, dependencies are also identified during big room planning and are tasks that require another Agile team to complete a different task(s) before action can be taken on the task planned by the first team. To address this challenge, dependencies are considered risks until a plan is established. If the team responsible cannot address the dependency in time, the dependency is considered an accepted risk (i.e., there is no mitigation strategy) or an issue in which an expected outcome is guaranteed to occur.

The Integrated Approach

Given the complex nature of short- and long-term risks, we suggest that the ideal RM system is one combining both approaches and in which the ROAM approach is prioritized for all risks and the traditional RM approach is used for long-term risks that cannot fit solely into the ROAM approach. This combined approach can be a game changer because it typically will require far less total investment to track and manage risks while providing more immediate feedback and insights.

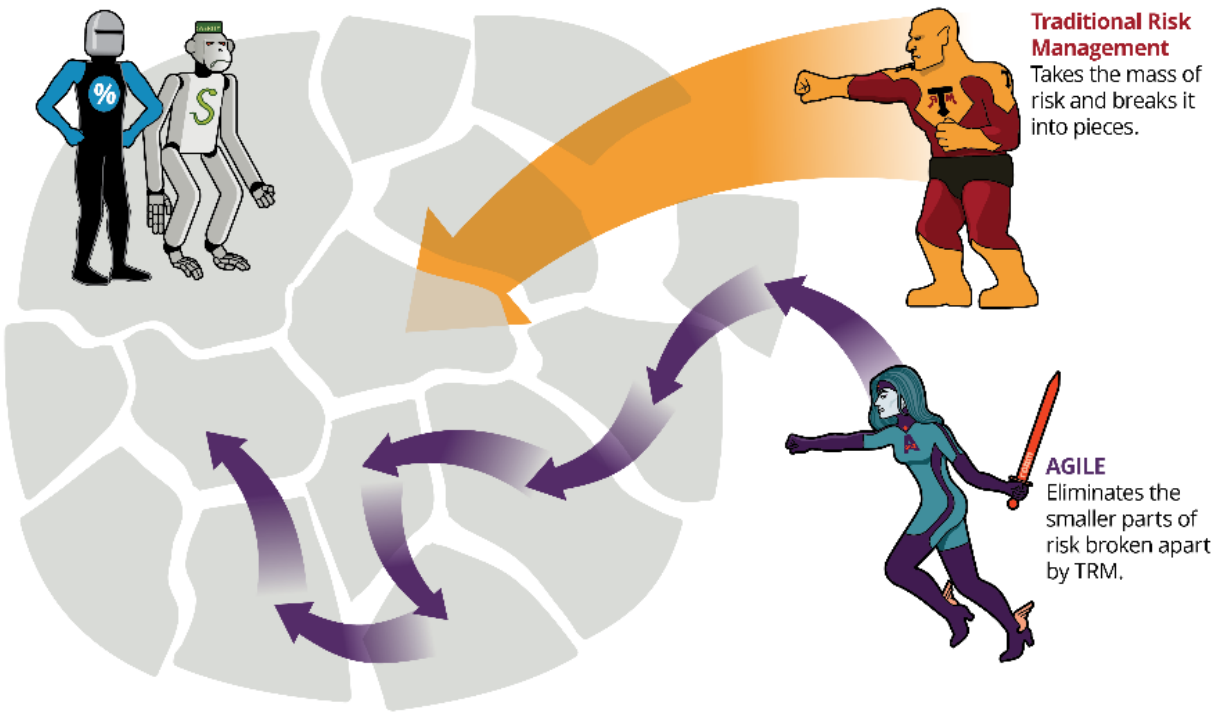

Figure 3. An RM approach combining traditional and Agile RM can cut costs, save time and effort, and lead to better, more efficient solutions.

To implement this combined approach in practice, we recommend that when a risk is identified, if possible, it should be evaluated as a short-term risk affecting the iteration work plan and introduced at the next ROAM event. It can and should be logged on a Risk ROAM board, but it does not need to be evaluated for criticality, reviewed at each periodic risk review board, and scrutinized for each mitigation step, as is typically the case with traditional risks.

To ensure an integrated approach is successful, the organization needs to continuously monitor existing risks. In this monitoring, the RM team needs to clearly communicate what is and is not a risk and how risks should be categorized and managed. This communication will reduce many common issues experienced by teams utilizing these processes. A common dashboard or tool with all the risks, traditional or ROAM, should also be utilized for a common risk repository. Having a dashboard or tool increases the likelihood that all risks are tracked appropriately because it acts as the single source of truth for all risks.

Additionally, clear and consistent lines of communication need to be established across risk managers, especially if the traditional RM work and the ROAM tasks are addressed separately. If this is the case, risk managers and stakeholders in the traditional RM approach often don't have a comprehensive view of all risks and some risks may get overlooked or underreported. Therefore, those established lines of communication create visibility that is essential in a combined approach.

In this integrated process, short-term risks will be addressed through the ROAM process as appropriate and long-term risks will not only be tracked but will also be better managed through the Agile-specific elements: incremental mitigation accomplished through the process of risk burndown, backlog grooming, and the creation of minimum viable products.

Risk Burndown

Figure 4. The probability and severity of identified risks can be significantly reduced when both traditional RM and Agile are used.

By using Agile with the traditional RM process, long-term risks will be captured and broken into smaller chunks that can be iterated on, mitigated, and easily tracked — a process called "burndown." In the combined approach, project managers, or "scrum masters" in Agile

terminology, will identify and manage the resolution of short-term RM activities related to long-term risks. They do this by first breaking down the long-term risks into bite-sized pieces and identifying which of those pieces are high priority and should be tackled during their planning cycle. The team then estimates how long it will take to accomplish each piece of work addressing the larger risk. The team can use this estimate to calculate how fast they can "burndown" or mitigate the risk by weighing the estimated effort against the estimated staffing efforts. Typically represented as a graph, this visualized estimate provides insight into how many risks exist in the planning cycle and how quickly the team will mitigate those selected risks. The team's progress is tracked and compared against the graph, providing easy reporting of the team capability and quicker identification and escalation of encountered challenges.

MVPs

An Agile approach also encourages experimentation, a process in which small prototypes, or minimum viable products (MVPs), are developed over a planning cycle — in accordance with the burndown chart — which provides stakeholders with more relevant and faster feedback. Using these MVPs can provide data for validating the risk mitigation hypothesis. By allowing for early and insightful feedback that can be implemented before changes are too expensive, this process of producing MVPs results in more trustworthy and meaningful risk mitigation.

Backlog Grooming

Because addressing long-term risks is work that's broken down into smaller chunks and prioritized in the Agile process, lower priority RM tasks may not be addressed within a given planning cycle. In the Agile process, these tasks are retained in the backlog, and the RM teams periodically review or "groom" the backlog to identify any tasks that are no longer needed, have had shifts in priority, or are connected to a high priority task. Tracking the smaller risk mitigation steps in an Agile project management tool in backlogs and Kanban boards allows the RM team and other stakeholders to visualize the smaller increments of work used to mitigate risks along with the other tasks that the Agile teams are working on, giving them a more comprehensive perspective.

Through this combined approach of breaking the large long-term risks into smaller more easily tracked risks, stakeholders are given better insight and ownership of tasks so that they know how to best manage resources and are given a clearer understanding of the program's risks. It also enables quicker mitigation or experimentation because, instead of waiting for a full solution based on a single hypothesis, teams can deliver faster, proven incremental solutions that gradually burndown the severity or probability from high to medium to low.

Why We Think a Combined Approach Will Work

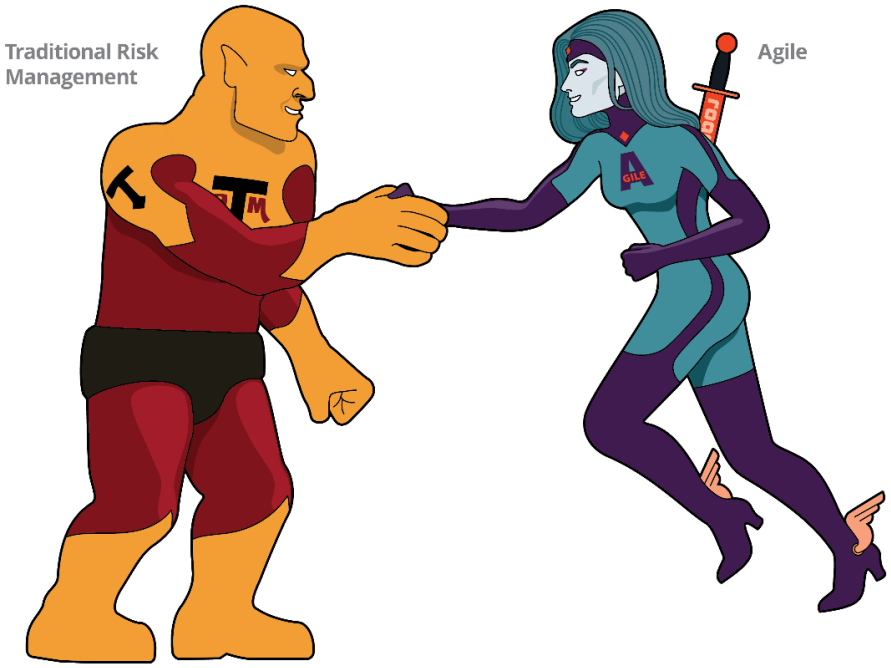

Figure 5. Despite taking different approaches, traditional RM and Agile RM have the same goals.

Interestingly, the RM approach outlined in the DoD's RIO Guidebook includes traditional RM guidance for programs *and* elements outlined in RIO are *analogous* to the Agile ROAM approach, revealing that the fundamentals of both approaches are aligned.

An example of the RIO Guidebook's analogous approach is seen with risk avoidance, which is when a program changes their plans to *avoid* or reduce the criticality of a risk. If an organization *transfers* a risk, they assign or share the risk with other entities or programs that are also impacted by that risk and may have more ability to mitigate it. In traditional RM, when an RM team seeks to *control* a risk, they are trying through whatever measures they have to reduce the probability of occurrence or the severity of the risk — essentially what all three of the RIO Guidebook mitigation options are trying to do. These processes are analogous to ROAM's mitigated state.

Therefore, the RIO Guidebook indicates that even though their processes differ, traditional RM and an Agile, lightweight approach like Risk ROAMing can work together because they have the same goals.

Conclusion

Organizations do not have the luxury of waiting months or years to fully close risks or reach an acceptable risk level; they must make incremental improvements that reduce severity or probability of occurrence in the most efficient way possible. Our proposed integrated RM approach helps organizations respond to ever-changing needs and provides a more insightful and less burdensome process. This combined approach saves time and money by requiring less reworking of long-term plans. Risk stakeholders and Agile teams also have better oversight of the work.

Breaking down the long-term risks allows for earlier experimentation to filter out ineffective alternatives and identify value. Additionally, with the increased visibility for the stakeholders, buy-in and understanding increases facilitation and allows for faster identification and categorization of the risks. Overall, this integrated approach empowers organizations to minimize their risks with more fidelity and efficiency, ensuring visibility of risks so that leadership can make the best decisions that maintain the mission, customer satisfaction, and trust for their programs.

References

[1] Himes, Emily. "ROAM & Risk Management Under SAFe®." PTC. 30 Oct. 2024. https://www.ptc.com/en/blogs/alm/roam-risk-management.

[2] Newman, Lily Hay. "Equifax Officially Has No Excuse." *WIRED*, 14 Sept. 2017, https://www.wired.com/story/equifax-breach-no-excuse/. Accessed 13 March 2025.

[3] U.S. Department of Defense. "Department of Defense Risk, Issue, and Opportunity (RIO) Management Guide for Defense Acquisition Programs." September 2023 - Incorporating Change 2.2 as of December 2023. https://www.cto.mil/wp-content/uploads/2024/05/RIO-2023-2-2.pdf.

[4] ISO. "ISO 31000:2018Risk management — Guidelines." Edition 2, 2018. https://www.iso.org/standard/65694.html#lifecycle.

Additional Reading

Place, Patrick R. and Stephen Wilson. "When Agile and Earned Value Management Collide: 7 Considerations for Successful Interaction." *SEI Podcasts* from Carnegie Mellon University Software Engineering Institute, 9 Feb. 2024, https://insights.sei.cmu.edu/library/when-agile-and-earned-value-management-collide-7-considerations-for-successful-interaction/.

Smith, Justin. "Incorporating Agile Principles into Independent Verification and Validation." *Carnegie Mellon University Software Engineering Institute*, 24 June 2024, https://insights.sei.cmu.edu/blog/incorporating-agile-principles-into-independent-verification-and-validation/. Accessed 13 March 2025.

Software Engineering Institute. "Agile Adoption in Government Podcast Series." *SEI Podcasts* from Carnegie Mellon University Software Engineering Institute, 27 Jan. 2017. https://insights.sei.cmu.edu/library/agile-adoption-in-government-podcast-series/.

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM25-0372