Category

Consulting

PMO & Transformation Governance Playbook (Status, Risks, Issues, RAID)

A practical method for running program governance using a disciplined RAID log, a consistent status cadence, and escalation thresholds that catch problems while they are still cheap to fix.

- Practitioner
- Intermediate
- Template Included

- Mithun A. Sridharan

- 28 April 2022
- 9 minute(s)

Overview

Run program governance with a disciplined RAID log, a consistent status cadence, and clear escalation thresholds, so risks get caught while cheap to fix instead of surfacing as expensive surprises.

What's the difference between a risk and an issue in a RAID log?

A risk is something that might happen in the future and would harm the program if it did; an issue is something that has already happened and is actively causing harm right now. The practical governance implication is that issues need faster review cycles and typically higher-urgency escalation than risks, and any risk that lingers unchanged for multiple review cycles should be scrutinized for whether it has quietly already become an issue.

How do you stop RAG status reporting from becoming subjective?

Define explicit, numeric thresholds before the program starts — for example, red means a milestone slip of more than two weeks or a budget variance beyond 10%, amber means a slip of one to two weeks or 5-10% variance — and apply them consistently across every workstream. Subjectivity creeps in whenever thresholds are vague ("significant delay") rather than numeric, because different leads calibrate "significant" differently.

Who should own an entry in the RAID log — the PMO or the workstream lead?

The workstream lead closest to the actual work should own risk, issue, and dependency entries relevant to their area, since they have the context to drive mitigation; the PMO's role

is to enforce the discipline — mandatory owners and dates, consistent RAG application, and flagging stale entries — rather than to personally own the mitigation actions themselves.

How do you handle a risk that nobody wants to escalate because it might reflect badly on their workstream?

Pre-agreed, objective escalation thresholds are the main defense — if a risk crosses a threshold defined before the program started, escalation is procedural rather than a judgment call the workstream lead has to personally initiate. Building a norm that a risk stuck at the same status for three reviews gets raised by the PMO automatically, not by the workstream lead volunteering it, also removes the incentive to quietly under-report.

How often should the full RAID log be reviewed versus just the steering committee summary?

Review the full RAID log with each workstream weekly (or biweekly for smaller programs) to catch stale items and drive mitigation action, and present a consolidated, threshold-based RAG summary to the steering committee on its own cadence — typically biweekly or monthly. Conflating the two into a single review either drowns the steering committee in workstream-level detail or leaves the working team without a regular forum to actually manage the risks.

{ "@context": "https://schema.org", "@type": "FAQPage", "mainEntity": [{ "@type": "Question", "name": "What's the difference between a risk and an issue in a RAID log?", "acceptedAnswer": { "@type": "Answer", "text": "A risk is something that might happen in the future and would harm the program if it did; an issue is something that has already happened and is actively causing harm right now. The practical governance implication is that issues need faster review cycles and typically higher-urgency escalation than risks, and any risk that lingers unchanged for multiple review cycles should be scrutinized for whether it has quietly already become an issue." } }, { "@type": "Question", "name": "How do you

stop RAG status reporting from becoming subjective?", "acceptedAnswer": { "@type": "Answer", "text": "Define explicit, numeric thresholds before the program starts — for example, red means a milestone slip of more than two weeks or a budget variance beyond 10%, amber means a slip of one to two weeks or 5-10% variance — and apply them consistently across every workstream. Subjectivity creeps in whenever thresholds are vague (\"significant delay\") rather than numeric, because different leads calibrate \"significant\" differently." } }, { "@type": "Question", "name": "Who should own an entry in the RAID log — the PMO or the workstream lead?", "acceptedAnswer": { "@type": "Answer", "text": "The workstream lead closest to the actual work should own risk, issue, and dependency entries relevant to their area, since they have the context to drive mitigation; the PMO's role is to enforce the discipline — mandatory owners and dates, consistent RAG application, and flagging stale entries — rather than to personally own the mitigation actions themselves." } }, { "@type": "Question", "name": "How do you handle a risk that nobody wants to escalate because it might reflect badly on their workstream?", "acceptedAnswer": { "@type": "Answer", "text": "Pre-agreed, objective escalation thresholds are the main defense — if a risk crosses a threshold defined before the program started, escalation is procedural rather than a judgment call the workstream lead has to personally initiate. Building a norm that a risk stuck at the same status for three reviews gets raised by the PMO automatically, not by the workstream lead volunteering it, also removes the incentive to quietly under-report." } }, { "@type": "Question", "name": "How often should the full RAID log be reviewed versus just the steering committee summary?", "acceptedAnswer": { "@type": "Answer", "text": "Review the full RAID log with each workstream weekly (or biweekly for smaller programs) to catch stale items and drive mitigation action, and present a consolidated, threshold-based RAG summary to the steering committee on its own cadence — typically biweekly or monthly. Conflating the two into a single review either drowns the steering committee in workstream-level detail or leaves the working team without a regular forum to actually manage the risks." } }] }

Subscriber access

Unlock this playbook

This playbook — including every framework, template, and step-by-step section — is available free to Think Insights subscribers. Enter your email to unlock it instantly and get our weekly insights newsletter. No account needed, and access is remembered on this device.

First Name

Last Name

Email

Country - Select -

■ I agree to data processing in accordance with GDPR

Submit

Leave this field blank

References

Related playbooks

Consulting

Trusted Advisor Storytelling & Thought Leadership Playbook

A framework for using genuine storytelling technique — specific narrative structure, concrete detail, and clear stakes — to make trusted advisor insight memorab

- Practitioner
- Intermediate
- Template Included

Read playbook

Consulting

Trusted Advisor Metrics & Self-Development Playbook

A framework for consultants to track their own trusted advisor development deliberately — structured feedback gathering, honest self-reflection, and targeted co

- Practitioner
- Intermediate
- Template Included

Read playbook

Consulting

Becoming a Trusted Advisor to Internal Stakeholders (In-House Consultant) Playbook

A framework for internal consultants, strategy teams, and in-house advisors to build genuine trusted advisor status with internal stakeholders — adapting extern

- Practitioner
- Intermediate
- Template Included

Read playbook

Talk to our consulting team

Author

Mithun A. Sridharan Thought Leader

I'm Mithun A. Sridharan, Founder of this website - Think Insights - on Strategy, Management Consulting, Leadership, Digital Transformation, and Data Literacy. Follow me on social media or connect with me on LinkedIn for updates.