Craft

GUIDE

The Federal Practitioner's Guide to Supply Chain Risk & Due Diligence

Frameworks, workflows, and decision tools for procurement, risk, and finance professionals — from threat identification through defensible award decisions.

CONTENTS

SECTION 01

Risk Ontology

Before you can assess risk, you need to understand what it is. Risk is not a number. It is not a score. It is the potential for loss or damage to something you value — and it only emerges when a hazard in objective reality threatens something valued under conditions of uncertainty.

> This matters for federal practitioners because risk tolerance varies legitimately across agencies. Army and Air Force organizations can look at the same vendor with the same risk profile and reach different decisions — not because one is wrong, but because their value frames differ. Understanding this prevents both over-rejection and under-screening.

SECTION 02

Evidence and Risk Effects

The chain from raw data to decision runs through evidence. Not all evidence works the same way — and misunderstanding the effect of evidence on risk is one of the most common sources of poor risk judgement.

Evidence effect	What it does to the risk structure	Practical effect
Reveals / Strengthens	Establishes or strengthens a hazard–exposure pathway	Increases the plausibility or consequence of the risk structure. Tightens the decision space.
Constrains / Weakens	Demonstrates barriers, controls, or separation between a hazard and the asset it might threaten	Limits plausibility or consequence. Can restore optionality for the decision-maker.
Clarifies	Improves understanding without changing plausibility	Reduces uncertainty about the risk structure without moving the posture recommendation one way or the other.

CORE PRINCIPLE

The analyst's job is evidence application, not data collection. Data becomes evidence only when it is linked to a risk finding. The question is never "what did we find?" — it is "what does what we found do to the risk structure, and therefore the decision space?" Analysts who skip this step produce reports that are hard to act on.

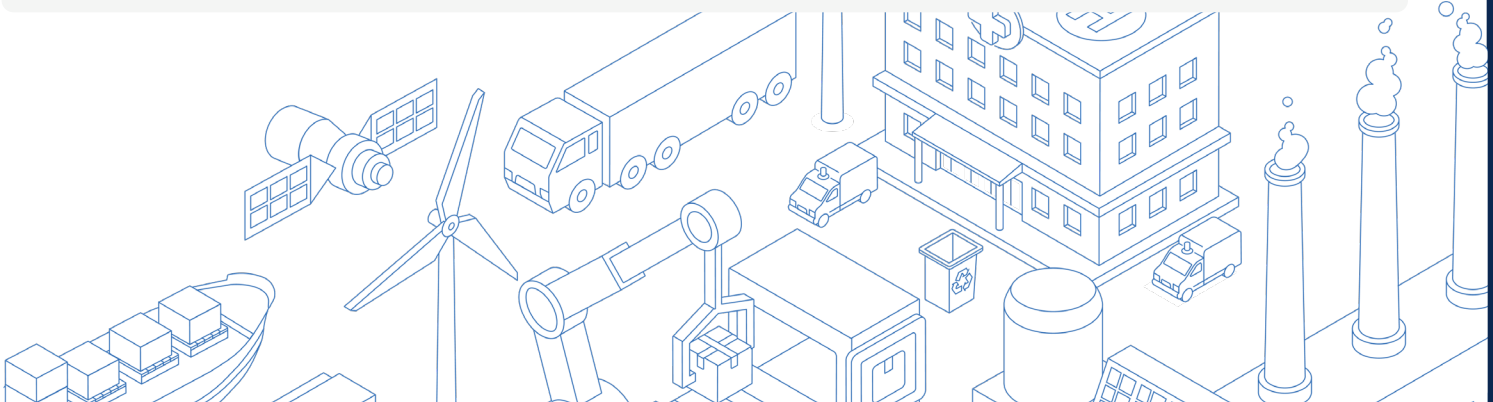

SECTION 03

Decision Geometry

The most actionable shift in due diligence thinking is moving away from low/medium/high risk scores toward classifying risk by its effect on the decision space. The analyst's core question is not "how risky is this?" — it is "what is the most reasonable and defensible posture given what we know?"

Decision-Constraining DC

Evidence is automatically linked to a prohibitive risk finding according to internal policy or law. No interpretation needed. These are structural triggers — typically "hard no" criteria (sanctioned entities, prohibited affiliations, banned parties).

Decision-Significant DS

Restricts the decision space without eliminating options. Requires deeper assessment or mitigation. A DS finding may shift to conditional acceptance if the company provides clarification or agrees to controls. DS can also be thought of as "notable" or "material", ultimately meaning risk which is relevant to the decision and impacts the decision space.

WORKED EXAMPLE: DECISION-SIGNIFICANT IN PRACTICE

A China-based private investor holds less than 20% equity in a vendor. This is a DS finding — not automatic rejection. The appropriate next step is company engagement to clarify whether the investor holds IP rights, board seats, or voting rights. If the company confirms none of the above and agrees to a silo investor condition, the posture shifts to Accept Conditionally. Reject remains on the table but is no longer the only defensible option. Risk assessment under uncertainty and in restricted decision space is where analyst assessment and judgement creates value.

> **Don't present a decision-maker with a low/medium/high score.**
> *Present them with a clear posture recommendation, the reasoning behind it, and the conditions under which it changes. That is a defensible decision brief. A risk score – number or letter grade - is not.*

SECTION 04

The FOCI/FCOC Threat

Foreign Ownership, Control, and Influence (FOCI) is the compendium of weapons adversaries use against federal supply chains. Below are examples of FOCI threat vectors, potential classifications of each, and mitigation levers. The level of risk is impacted by which foreign country is involved. Involvement of a foreign country-of-concern (FCOC) amplifies risk significantly.

Threat vector	Key indicators	Default geometry	Mitigation lever
Sanctioned / denied-party affiliation	OFAC/BIS hits; prohibited entity list matches; unlicensed dual-use exports	DC — Reject	No mitigation — structural trigger
Foreign majority / board control	FCOC majority equity; FCOC C-suite appointment rights; veto clauses	DC / DS	Government Security Agreement; proxy board; Security Control Agreement
Civil-military fusion affiliation	Subsidiaries in CMF countries (FCOC, especially); shared IP with state-linked entities; joint R&D	DC / DS	Corporate-wide Technology Control Plan; divestiture requirement
Insider threat exposure	Personnel with undisclosed FCOC contacts; access gaps; compliance failures	DS	Enhanced foreign travel/contact reporting; access tiering; personnel restrictions
Minority FCOC investment (<20%)	FCOC PE/VC stake; opaque holding structures; country-of-concern origin	DS	Silo investor condition; no IP / board / voting rights confirmation
Supply chain dependency	Sole-source foreign FCOC suppliers for critical components; restricted-mineral exposure	DS	Diversification requirements; domestic sourcing; buffer stock
Cyber vulnerability	FCOC -developed software; unpatched infrastructure; FCOC cloud dependencies	DS	CMMC compliance; network segmentation; vendor cyber attestation
Market denial (commercial)	Post-acquisition exit from U.S. market; FCOC buyer acquires niche domestic capability	DS	CFIUS review; domestic capability investment; alternative supplier development

Decision-constraining findings make decisions easy. If the findings meet hard denial criteria and the risk is non-mitigable, little thought is necessary. Within decision-significance, risk assessment is complicated. It is not practically useful to attach a label or risk rating within decision-significance. Each risk instance is unique. Compressing risk complexity with a number or letter grade only leads to confusion. The only time a weighted risk score may be valuable is during triage and prioritized execution of assessment.

> *But it MUST not be confused with a legitimate, final risk assessment and recommended decision posture. A decision-significant finding simply means that risk is notable enough to bring to a decision board.*

SECTION 05

Scaled Due Diligence: The End-to-End Workflow

Establish or license a data infrastructure with entity resolution across beneficial ownership, corporate registries, sanctions lists, denied-party databases, and open-source information. Data quality and coverage at this stage determines the confidence of every downstream decision.

1

Triage — clearing and escalation at scale

Apply pre-defined risk signal thresholds to the full company population. Companies that breach no thresholds are mechanically cleared (Decision-Neutral). Companies that breach thresholds are escalated to analyst review. The goal is to protect analyst capacity for cases where human judgement adds value — not to have analysts manually review every company in a 17,000-proposal pipeline. In addition to thresholds for individual data points, there may be analytic or AI-driven thresholds for flagging companies.

2

Analysis & assessment – human analyst + AI support

For escalated companies, analysts examine risk structure: identify hazards, map exposure pathways, evaluate controls, and classify each finding as DC (hard no), DS (notable/material), or DN (reviewed but insignificant risk/not relevant to decision). AI-assisted research compresses the time to synthesize corporate documentation and open-source data. The analyst's job is not to collect more data — it is to link evidence to risk findings, assess the overall risk structure, and judge the most reasonable and defensible posture.

3

Consortium model: the efficiency multiplier

OCEA data shows 60–70% overlap in companies applying for awards across Space Force, Air Force, and NASA. Independent screening at each agency means paying three times for the same analysis — and often reaching inconsistent conclusions. The consortium model (shared platform, shared findings, agency-specific award decisions) reduces cost, improves consistency, and preserves agency autonomy in the final decision.

SECTION 06

The Defensible Decision Brief: A Template

DECISION-SIGNIFICANT FINDINGS

Only include findings that materially shape the reasoning.

- Linked risk findings
- Exposure structure
- Why it requires consideration
- Whether controls meaningfully offset it

CONDITIONS (IF APPLICABLE)

Only include if posture is conditional. Each condition maps to a risk finding.

- Required clarification from company
- Required mitigation before award
- Ongoing compliance requirements post-award

APPENDIX (AUDITABILITY ONLY)

Do not include it in the oral brief. Exists for audit trail and deep dive.

- AI risk analysis output
- Analyst in-depth assessment
- Source documentation

SECTION 07

The Program Manager's Pre-Award Checklist

For PMs who judge risk, the challenge is interpreting it correctly and asking the right questions before committing to a source selection decision.

PRE-AWARD CHECKLIST

- [ ] Have I identified whether any findings are Decision-Constraining? If yes, the decision is made — document and move on.
- [ ] Do I understand which specific risk findings are Decision-Significant, and what each means for my program's specific technology and classification level?
- [ ] Have I mapped the vendor's risk profile against the technology sensitivity of this contract — not a generic program risk profile?
- [ ] For DS findings: have I asked what conditions or mitigations would shift the posture to Accept Conditionally? Is the vendor willing to engage?
- [ ] Have I checked whether this vendor has been screened by other agencies — and what posture those agencies adopted?
- [ ] For supply chain contracts: have I requested sub-tier analysis beyond the prime? Foreign exposure at tier 2 or tier 3 can be a real risk.
- [ ] Is a continuous monitoring requirement in place post-award? Who owns it, and what triggers re-assessment?
- [ ] Have I documented my risk rationale — including the posture I selected and why — regardless of whether I award or deny?
- [ ] If this is a re-compete or follow-on: have I re-screened the vendor? Prior awards carry no current-cycle clearance.
- [ ] Do I know which materials in this supply chain have foreign-sourcing exposure (gallium, germanium, antimony, and five additional restricted materials)?

SECTION 08

Field-Tested Lessons Learned

These are practitioner-tested principles from federal risk analysts working at scale. They apply regardless of agency, program type, or contract value.

- **Understand risk conceptually before you assess it.** Analysts who treat risk as a score rather than a structure consistently produce reports that decision-makers cannot act on.
- **Leave room for nuance and context — it always matters.** A “decision-significant” vendor making rucksacks is not the same as a “decision-significant” vendor making hypersonic components.
- **Don’t brief low / medium / high.** Brief the defensibility of the decision posture. “The most reasonable and defensible recommendation is Accept Conditionally” is actionable. “Medium risk” is not.
- **Triage, prioritize, execute.** Protect analyst capacity for Decision-Significant cases. Do not let the volume of Decision-Neutral companies crowd out the ones that actually need human attention.
- **Clarify and document “hard no” criteria before you start screening.** DC triggers should be defined and agreed upon with decision-makers in advance, not discovered mid-assessment. That said, DC criteria will evolve with time.
- **Prioritize in order: DC → DS → DN.** Work the constrained cases first. Decision-Neutral findings are closure, not analysis tasks.
- **Risk avoidance is impossible.** The goal is always mitigation to an acceptable level — not elimination. Rejection is appropriate when mitigation cannot reduce risk to an acceptable posture, not as a reflex to any flag.
- **Know when to stop the deep dive.** At a certain point, the risk picture is clear enough to judge and inform a decision. More data does not always improve the recommendation. Recognize that point and stop.
- **In briefing: conciseness, clarity, confidence, sound reasoning.** Decision-makers are not risk analysts. Your job is to translate the analysis into a posture recommendation they can own. Lead with the BLUF, not the methodology.