COURSE SYLLABUS

FedRAMP for Agency Buyers: Package Due Diligence and Cloud Security Program Management (Advanced) Course

Take FedRAMP package due diligence to depth, hold CSPs to account on continuous monitoring obligations, and design enterprise governance for multi-cloud federal portfolios.

Group classes in Live Online and onsite training is available for this course. For more information, email onsite@graduateschool.edu or visit: https://www.graduateschool.edu/courses/fedramp-for-agency-buyers-package-due-diligence-and-cloud-security-program-management-advanced-course

support@graduateschool.edu •
(888) 744-4723

Course Outline

Module 1: SSP Architecture and Authorization Boundary Analysis

- Decomposing the authorization boundary: service layers, management plane, and supporting infrastructure the CSP places inside versus outside the boundary.
- Data flow diagram analysis: federal data ingress and egress paths, metadata and telemetry flows, and where diagrams contradict the system narrative.
- Component inventory reconciliation: verifying the inventory agrees with the boundary diagram and with the environment actually assessed.
- External services, corporate systems, and subservice organizations: what sits outside the boundary, what risk transfers to the agency, and how leveraged authorizations are disclosed.
- Investigating boundary ambiguity: unstated dependencies, undocumented management interfaces, CSP-side administrative access paths, and the questions that force each one to resolution.
- Architecture implications for the agency: which CSP components the agency must represent in its own authorization boundary.

Module 2: SSP Control Implementation Analysis and Inheritance Verification

- Organization-defined parameters: how CSP ODP selections constrain agency flexibility and where they fall short of agency policy.
- Control origination analysis: verifying claimed inheritance against what the leveraged provider's package actually supports.

- Inheritance chain testing: obtaining the leveraged provider's package and proving the control is implemented at the tier that claims it rather than assumed at every tier.
- CRM sufficiency testing: identifying customer-responsible controls described too vaguely to implement or evidence, and the specificity the agency should demand in writing.
- Cross-referencing SSP claims against the SAR and POA&M to surface internal contradictions in the package.

Module 3: SAR Analysis — Assessment Scope, Evidence Quality, and Finding Calibration

- Adjudicating evidentiary sufficiency: distinguishing a control the assessor tested weakly from one never tested at all, deciding how much the agency may infer from sampling depth and assessor judgment, and defending that threshold when the CSP disputes it.
- Assessment scope testing: comparing tested components against the boundary and inventory to expose elements never assessed.
- Sampling adequacy: population definition, sample size, and when a sample cannot support the conclusion drawn from it.
- Finding severity calibration: distinguishing material from administrative findings and building the case to challenge an under-rated finding.
- Remediation commitments in the SAR: assessing credibility and verifying they carry through to the POA&M.

Module 4: Assessor Qualification, Package Currency, and CSP Evidence Demands

- 3PAO accreditation and assessment team qualifications: verifying the assessor and the individuals who performed the work.
- Assessor independence and conflict of interest: advisory-to-assessment conflicts, repeat-engagement risk, and CSP influence over scope.
- When accumulated change outruns the evidence: reasoning about the gap between the system as assessed and the system as it now runs, what can be inferred about that gap from monitoring deliverables, and when no authorization decision can defensibly rest on the package.
- Formal evidence requests: what an agency can demand from a CSP, how to frame the request, and escalation paths when the CSP declines.
- Documenting due-diligence limitations: recording what could not be verified and carrying it forward into residual risk.

Module 5: Container Security Compliance in FedRAMP Environments

- FedRAMP container scanning requirements: hardened image criteria, registry scanning obligations, and the image rebuild and redeployment cadence the CSP is required to meet.
- Containers in the authorization boundary: how containerized services appear in the CSP's inventory, what the agency inherits, and how a boundary is drawn around a workload set whose composition changes continuously.
- Image hardening and provenance: approved base images, baseline configuration compliance, signing and attestation, and NIST SP 800-190 alignment.
- Orchestration platform controls: control-plane hardening, namespace and tenant isolation, secrets management, and admission control.
- Container evidence in the SSP and SAR: verifying the 3PAO assessed the container and orchestration stack rather than only the host layer.

- Reviewing the monthly container deliverable: reconciling ephemeral and short-lived assets against the submitted inventory, and confirming image vulnerabilities are carried and aged correctly on the POA&M.

- Agency-side responsibilities: what the CRM leaves to the agency when agency teams deploy their own images and workloads onto an authorized platform.

Module 6: AI/ML Cloud Services — Authorization Gaps and Agency Safeguards

- FedRAMP status of AI/ML services: authorization gaps, and what “authorized” actually covers when a CSP layers AI features onto an existing service.

- AI features as significant changes: SCN triggers when a CSP adds generative or ML capability to an authorized service, and the agency's authorization impact review.

- OMB M-24-10 and federal AI policy intersection: high-impact AI use cases, minimum practices, and how they layer on top of a FedRAMP ATO.

- AI-specific risk areas beyond the Rev 5 baseline: training data handling, prompt and inference logging, model provenance, tenant data isolation, and retention for model improvement.

- Where 800-53 Rev 5 falls short for AI workloads: identifying control gaps and selecting supplemental agency controls (NIST AI RMF as a supplement, not a substitute).

- Contractual and CRM implications: data-use restrictions, training opt-out, output rights, model change notification, and the evidence an agency should demand.

Module 7: Multi-Cloud Authorization Architecture and Control Inheritance

- Boundary decomposition across providers: deciding where the agency boundary cuts through each provider's service when three CSPs each claim a different edge, and documenting a cut that survives assessor challenge.

- Conflicting inheritance across providers: reconciling different customer-responsibility assertions for the same underlying platform.

- CSP-to-CSP interconnections: assessing provider-to-provider data flows the agency does not control, and establishing agency visibility and notification rights over them.

- When evidence from two regimes is not comparable: a sample-based 3PAO control test and an RMF assessment of the same control family yield different assurance — normalizing sampling depth, test date, and assessor independence before the AO sees one risk picture.

- Authorization strategy for in-process, lapsed, and exiting services: relying provisionally on a package still under review, tracking a CSP whose authorization has expired or been withdrawn, and the agency's fallback position when a service leaves the Marketplace mid-contract.

- The carrying cost of a boundary decision: how a consolidated versus federated authorization structure changes recurring monitoring workload, staggers assessment calendars against each other, and determines whether one component's failure suspends the whole system or a severable piece of it.

Module 8: Enterprise Cloud Portfolio Governance and Vendor Accountability

- Portfolio inventory: authorizations, expiration and reauthorization dates, significant change status, and responsible ISSOs.

- Risk aggregation across CSPs: common vulnerabilities, shared upstream providers, and concentration risk from a single point of failure.

- Longitudinal vendor scoring: tracking each CSP's performance trend over time and comparing providers against one another across the portfolio, rather than reading the same signals once at selection.
- Enforcement leverage the agency actually holds: separating rights written into the contract vehicle from expectations the agency merely asserts, and building enforceable security obligations in before award.
- Portfolio-level decision-making: consolidation, exit, and reauthorization prioritization under resource constraints.

Module 9: Agency ATO Documentation for FedRAMP-Leveraged Systems

- Assembling the leveraged authorization package: sequencing the memo, supplemental documentation, and evidence set so each artifact substantiates the next rather than restating it.
- Documenting contested boundary placement: recording the rationale for CSP management planes, shared services, and agency-configured components so an assessor can follow the decision without re-litigating it.
- Package mechanics for containerized and AI-enabled components: where those components are represented across the memo, supplemental SSP, and boundary artifacts so the package stays internally consistent.
- ISA sufficiency for cloud interconnections: adapting agreements built for system-to-system connections to CSP relationships where the agency cannot dictate terms.
- Package traceability: making inherited, shared, and agency-owned control coverage provable end to end within the agency package.

Module 10: Residual Risk Analysis and the Authorization Decision

- Residual risk derivation from three sources: open SAR findings, active POA&M items, and due-diligence items the agency could not verify.
- Aggregating residual risk: converting individual findings into a system-level risk position the AO can weigh, expressed in mission likelihood and impact rather than control language, without collapsing it into a single score.
- Risk accepted on behalf of others: quantifying what a downstream agency inherits when it reuses this authorization, and what the agency owes that reuser in disclosure.
- Engineering enforceable conditions: pairing each condition with its verification evidence, deadline, responsible party, and consequence for non-performance.
- Conditions the agency cannot enforce: determining whether privity, a reseller intermediary, or a shared vehicle leaves the AO any mechanism standing behind a condition, and what changes in the decision itself — narrower scope, shorter term, or declining to authorize — once a condition is known to be unbacked.
- Defending the decision under later review: building the record that supports the authorization when an Inspector General, auditor, or successor AO reexamines it.

Module 11: Capstone — Complete Agency ATO Package Development

- Capstone scenario: a candidate FedRAMP Moderate service running on a containerized platform, with a recently added AI capability and dependencies on two authorized services already in the agency portfolio.
- Executing package assembly under deadline: producing the full artifact set — due-diligence memo, supplemental SSP, boundary and data flow diagrams, ISA, residual risk statement, conditions of authorization, and continuous monitoring plan — against the sequencing discipline established earlier in the course.

- Applied quality review: running another team's completed package against the Inspector General defensibility standard and returning the findings.
- The authorization decision: presenting the package to the Authorizing Official, defending boundary and inheritance assumptions, and adjudicating challenges to residual risk.
- Post-authorization handoff: transferring the package, the conditions verification schedule, and CSP deliverable tracking to the ISSO who will carry it.
- After-action review: translating capstone outcomes into improvements to the agency's standing authorization process.