ISC2™

ISC2 **Exam Guidance** for Artificial Intelligence

April 2, 2026

ISC2 Exams and AI Security Concepts

As artificial intelligence (AI) rapidly transforms how organizations operate, cybersecurity professionals face a new era of risk and opportunity. Threat actors are leveraging AI to discover vulnerabilities, accelerate reconnaissance and launch highly personalized social engineering attacks. At the same time, AI is helping cybersecurity professionals automate repetitive tasks, improve threat detection, increase efficiency and reduce human error.

The ISC2 Exam Guidance for Artificial Intelligence outlines how AI-related knowledge, skills and abilities are incorporated across ISC2's portfolio of globally recognized vendor-neutral cybersecurity certification exams. This document maps where AI concepts appear within more than 50 core cybersecurity exam domains and 200 subtasks, demonstrating how ISC2's rigorous and continuous exam maintenance process is ensuring that certified professionals remain prepared to secure organizational assets in an increasingly AI-driven world.

ISC2 Certification Exams Continuously Evolve with Emerging Practices

Through a rigorous, 3-year exam refresh cycle—including Job Task Analysis (JTA), exam blueprint development, item writing, peer review, standard setting and publishing—our certified Subject Matter Experts and practitioners from the field ensure that ISC2 exams reflect real-world professional requirements. As AI capabilities evolve and intersect core cybersecurity domains, these experts regularly integrate AI-focused tasks and security considerations into the certification exam blueprints, ensuring that ISC2 credentials remain relevant, timely and rigorous.

AI security concepts have been continually integrated throughout ISC2's core cybersecurity domains, spanning Security and Risk Management, Asset Security, Security Architecture and Engineering, Communication and Network Security, Security Assessment and Testing, Security Operations and Software Development Security and more.

Explore the ISC2 Certification Exam Domains that Incorporate AI Security

Mapped to all **nine certification exams in the ISC2 portfolio**, this Exam Guidance ensures cybersecurity professionals and their employers can be confident that all ISC2 accredited certifications remain on the leading edge of today's AI cybersecurity practices.

Table of Contents

Page 3	CC
Page 5	SSCP
Page 8	CISSP
Page 11	CCSP
Page 14	CGRC
Page 17	CSSLP
Page 20	ISSAP
Page 22	ISSEP
Page 24	ISSMP

Based on CC Exam Outline Effective September 1, 2026

The Certified in Cybersecurity (CC) certification is the global gateway for individuals entering the cybersecurity workforce. As AI becomes a standard component of corporate technology, it is essential that even entry-level professionals understand its security implications. For the CC Exam Outline, we have integrated foundational AI concepts across all five domains. This approach ensures that new practitioners can identify AI assets, recognize automated threats and support the governance frameworks that keep these emerging technologies secure.

Domain 1: Security Principles

At the foundational level, the CC Exam Outline introduces how AI impacts the core pillars of information security: Confidentiality, Integrity and Availability. Entry-level professionals understand how to apply the fundamental security principles to AI systems, specifically focusing on how data integrity is vital for preventing "model poisoning." The integration also covers the ethical use of AI, highlighting the importance of transparency and non-bias in automated decision-making as part of a comprehensive security culture.

Furthermore, this domain establishes the role of AI within the broader governance, risk and compliance (GRC) landscape. Candidates are aware that AI tools are subject to the same organizational policies and legal requirements as traditional software. By understanding these high-level principles, new practitioners can support senior leadership in ensuring that AI adoption aligns with the organization's risk appetite and ethical standards.

Domain 2: Business Continuity (BC), Disaster Recovery (DR) & Incident Response (IR) Concepts

In this domain, the CC Exam Outline incorporates how AI both complicates and enhances the resilience of an organization. From a response perspective, candidates understand the basics of how AI-driven tools can assist in the early detection of security incidents. The CC Exam Outline emphasizes the role of the entry-level practitioner in following established playbooks that now account for automated threats, ensuring they can provide valuable support during the initial triage of a suspected breach.

Regarding recovery and continuity, the integration focuses on the necessity of backing up not just traditional data, but the specific configurations and datasets that power AI services. The CC Exam Outline includes the concept of "Model Drift" as a potential business continuity risk, where an AI's declining performance could impact critical operations. By using these concepts, CC holders are prepared to assist in maintaining the availability of intelligent systems during and after a disruptive event.

Domain 3: Access Control Concepts

Access control is the first line of defense. Candidates understand that just like human users, AI "bots" and automated service accounts must be managed through a formal lifecycle—from provisioning to deprovisioning. The integration emphasizes the Principle of Least Privilege, allowing entry-level staff to verify that automated systems only have the permissions necessary to perform their designated tasks.

Additionally, the CC Exam Outline includes how AI is used to strengthen authentication through behavioral analysis. Candidates understand the foundational concepts of Multi-Factor Authentication (MFA) and how AI can help detect "impossible travel" or anomalous login patterns. This ensures that new professionals understand both how to secure the AI's access and how AI serves as a silent partner in protecting user identities across the enterprise.

Domain 4: Network Security

For network security, the CC Exam Outline incorporates the basics of how AI influences traffic monitoring and threat prevention. Entry-level practitioners understand the concept of AI-powered firewalls and Intrusion Detection Systems (IDS) that go beyond simple signature matching. By understanding how these tools use machine learning to identify unusual network behavior, candidates are better equipped to monitor dashboards and report potential anomalies to senior analysts.

The integration also addresses the security of the pathways that AI data travels. The CC Exam Outline describes the importance of network segmentation to keep AI development environments isolated from sensitive production data. This foundational knowledge allows CC professionals to support the implementation of Zero Trust principles, ensuring that the network remains a secure environment for the transmission of high-value AI training data.

Domain 5: Security Operations

In the final domain, the CC Exam Outline focuses on the day-to-day tasks of a security professional working alongside AI. This includes the foundational understanding of how Security Information and Event Management (SIEM) tools use AI to correlate data and reduce "alert fatigue." Candidates know how to handle the outputs of these automated systems, ensuring they can distinguish between a routine automated block and a high-priority event that requires human intervention.

We also introduce the "Security of the AI Workspace," focusing on the safe use of web-based AI tools and LLMs in an office environment. Candidates are prepared to identify the risks of "Data Leakage" when employees interact with public AI services. Subtasks within this domain ensure that candidates can serve as effective "human firewalls," protecting the organization's data integrity in an increasingly automated world.

Based on SSCP Exam Outline Effective September 15, 2024

The Systems Security Certified Practitioner (SSCP) remains the premier certification for hands-on security administrators. As AI tools transition from experimental to operational, the SSCP Exam Outline has evolved to ensure practitioners can implement, monitor and administer these technologies safely. By embedding AI topics into the seven SSCP domains, candidates are tested to manage the technical realities of AI—from securing automated access controls to leveraging machine learning for real-time incident response.

Domain 1: Security Concepts and Practices

In the foundational domain of the SSCP, AI integration centers on the fundamental shift in how we apply the pillars of information security to automated systems. Security administrators must now understand how "Algorithmic Integrity" ensures that AI outputs remain reliable and untampered. We incorporate AI by emphasizing the importance of ethical guidelines and transparency in automated processes, ensuring that as systems become more autonomous, they remain accountable to established security policies and organizational standards.

Furthermore, this domain addresses the lifecycle management of AI-enabled security controls. As a practitioner, you are tasked with supporting change management processes that account for the unique update cycles of Machine Learning (ML) models. This integration ensures that the security professional is not just a passive observer but an active participant in maintaining the functional security posture of AI-driven business tools.

Domain 2: Access Controls

Access control in the age of AI requires managing a new class of "intelligent" non-human entities. Within this domain, we focus on the identity management lifecycle for AI agents and service accounts that perform automated tasks. The integration covers the implementation of the Principle of Least Privilege for these agents, preventing them from accessing sensitive data tiers that fall outside their specific operational scope or training requirements.

Additionally, we include how AI enhances traditional access control mechanisms. By supporting the implementation of adaptive authentication and behavioral biometrics, SSCPs can utilize AI to detect anomalies in user access patterns in real-time. This dual-sided integration ensures that while you are securing the AI's access to your systems, you are also leveraging AI to make your overall access control architecture more resilient and dynamic.

Domain 3: Risk Identification, Monitoring and Analysis

This domain shifts the practitioner's focus toward the visibility and reporting of AI-specific risks. Organizations can integrate AI by training administrators to identify Indicators of Compromise (IoC) that are unique to ML environments, such as "Model Drift" or suspicious query patterns. Candidates have knowledge on security assessments that evaluate the vulnerability of AI endpoints, ensuring that these systems are captured within the organization's broader risk register and vulnerability management lifecycle.

Monitoring also takes a significant leap forward with the inclusion of AI-driven analytics. SSCPs are tasked with analyzing results from correlation engines that use machine learning to reduce noise and highlight true security events. By incorporating these automated visualization and trend-analysis tools, practitioners can communicate findings more effectively and escalate critical risks before they result in a breach.

Domain 4: Incident Response and Recovery

When a security incident occurs, speed is the most critical factor. In Domain 4, we integrate AI by focusing on the use of automated playbooks and AI-assisted triage during the initial response phase. Security practitioners can support forensic investigations where AI may have been either the target or the perpetrator, ensuring that evidence handling—such as the "chain of custody" for model logs—is conducted according to legal and ethical principles.

Recovery operations now account for the unique requirements of restoring AI-driven systems. This includes ensuring that backup and recovery procedures for ML models and their associated training data are robust. By integrating AI into incident response, we ensure that the modern security administrator can maintain business continuity even in environments where threats and defenses move at machine speed.

Domain 5: Cryptography

The cryptography domain addresses the essential role of encryption in protecting datasets that fuel AI. Integration in this domain focuses on securing "Data in Use" during the training and inference phases, utilizing modern cryptographic protocols to prevent data leakage. We also test the implications of quantum computing and advanced cryptanalysis on the long-term security of AI assets, ensuring that the practitioner can implement resilient key management for AI-related secrets.

Furthermore, we address the use of blockchain and other distributed ledger technologies to provide non-repudiation for AI decision-making. By applying cryptographic signatures to the outputs of AI models, SSCPs can help ensure that the provenance of an automated decision is verifiable and that the data has not been modified in transit, maintaining the trust required for autonomous operations.

Domain 6: Network and Communications Security

As AI workloads scale, they place unique demands on network architecture. This domain integrates AI by focusing on the secure placement and configuration of network-based security devices that monitor AI traffic. Practitioners have knowledge of implementing micro-segmentation to isolate AI training clusters, preventing an adversary from using a compromised AI interface as a beachhead to move laterally across the organizational network.

We also cover the role of AI in defending the network perimeter. This includes the administration of AI-powered firewalls and intrusion prevention systems (IPS) that can detect and block sophisticated "Low and Slow" attacks. By securing the communication pathways for Internet of Things (IoT) devices and mobile endpoints that utilize AI, the SSCP ensures that the network remains a hardened environment for intelligent applications.

Domain 7: Systems and Application Security

In the final domain, the SSCP Exam Outline focuses on the day-to-day administration of the systems that host AI applications. Integration involves managing the software supply chain for ML libraries and ensuring that application security testing accounts for AI-specific logic flaws. Practitioners are tasked with overseeing the secure deployment and patching of these systems, ensuring that "Model Hijacking" or inference attacks are mitigated at the application layer.

Mobile and cloud security also play a major role, as many AI services are delivered via these platforms. SSCPs have knowledge of administering containerized AI environments and manage the security of the APIs that connect these applications to the rest of the enterprise. This ensures that from the server to the endpoint, the entire application stack is resilient against both traditional vulnerabilities and modern, AI-augmented threats.

Based on **CISSP Exam Outline** Effective April 15, 2024

As AI and machine learning (ML) become foundational to modern business operations, the CISSP certification has evolved to ensure that cybersecurity professionals can govern, design and defend these sophisticated systems. Rather than treating AI as a siloed topic, ISC2 continues to interweave AI-specific security tasks and subtasks across all eight domains of the CISSP Exam Outline. This ensures a holistic approach to security that addresses the unique risks of algorithmic bias, data poisoning and adversarial attacks while leveraging AI for defensive automation.

Domain 1: Security and Risk Management

Security leadership now requires a deep understanding of how AI assets shift the organizational risk posture. Within this domain, the CISSP Exam Outline emphasizes the integration of ML models and LLMs into existing risk management frameworks. This includes establishing governance for AI ethics and mitigating algorithmic bias, ensuring that automated decision-making processes align with legal, regulatory and privacy requirements.

Furthermore, AI integration touches upon third-party risk management. As organizations increasingly rely on external AI service providers, CISSPs must be equipped to evaluate the security of AI supply chains. This involves assessing the transparency of data sourcing and the resilience of provider-managed models against evolving threats, ensuring that AI adoption does not create unmanaged blind spots in the corporate security strategy.

Domain 2: Asset Security

In the realm of asset security, data is the lifeblood of AI, and its protection is paramount. This domain now covers the classification and handling of AI-specific assets, such as training datasets, pre-trained models and model weights. We focus on maintaining data integrity throughout the AI lifecycle, ensuring that the information used to “teach” these systems has not been tampered with or poisoned by malicious actors.

Privacy remains a cornerstone of this domain, specifically regarding how AI systems process Personally Identifiable Information (PII). Integration efforts focus on technical controls like differential privacy and data masking within AI environments. By treating ML models as high-value intellectual property, we provide a roadmap for managing the collection, storage and eventual destruction of data in a way that satisfies both security and privacy mandates.

Domain 3: Security Architecture and Engineering

The architecture and engineering domain addresses the structural defenses required to host and run AI safely. This includes the design of secure enclaves for high-performance AI compute and the implementation of robust input-validation mechanisms to defend against prompt injection and adversarial attacks. The CISSP Exam Outline integrates AI by looking at the shared responsibility models inherent in cloud-based AI services, ensuring the underlying infrastructure is resilient to the unique computational demands of neural networks.

Beyond physical and logical hosting, this domain includes the engineering of "Explainable AI" as a security requirement. By building systems that provide transparency into how they reached a specific output, security engineers can better audit AI behavior. This integration ensures that security architecture isn't just a perimeter around a system, but a transparent framework that supports the verification and validation of AI-driven security controls.

Domain 4: Communication and Network Security

As AI workloads move across the network, this domain focuses on securing the transit of massive datasets and the communication between distributed AI nodes. Integration involves implementing specialized micro-segmentation and Zero Trust Architecture (ZTA) to isolate AI training environments from the rest of the enterprise network. This prevents lateral movement in the event of a compromised AI interface.

Additionally, we address the role of AI in network defense. CISSPs are tasked with understanding how AI-driven Network Detection and Response (NDR) tools identify anomalous traffic patterns that traditional signature-based systems might miss. By securing the channels used for "inference at the edge," we ensure that the communication pathways supporting AI remain confidential and available.

Domain 5: Identity and Access Management (IAM)

Identity remains the primary perimeter in an AI-driven world. Within Domain 5, the CISSP Exam Outline focuses on managing identities for non-human entities, specifically AI agents and automated service accounts. This integration ensures that AI systems operate under the Principle of Least Privilege, preventing "privilege escalation" where an AI might gain unauthorized access to sensitive data repositories during its learning or execution phase.

The CISSP Exam Outline also incorporates the use of AI to enhance IAM through behavioral biometrics and adaptive authentication. By leveraging AI to analyze user login patterns and detect anomalies in real-time, CISSPs can implement more dynamic access controls. This dual focus ensures that while we secure the AI's identity, we also use AI to make the entire organization's identity infrastructure more resilient.

Domain 6: Security Assessment and Testing

Security testing must now evolve to include "Red Teaming" for AI systems. Within this domain, the CISSP Exam Outline integrates methodologies for testing model robustness against evasion and extraction attacks. Professionals audit AI systems not just for software bugs, but for "logic flaws" in the model's output that could be exploited by an adversary.

Furthermore, we address the use of AI to automate the vulnerability management lifecycle. By integrating AI-powered scanning tools, organizations can prioritize remediation efforts based on real-time threat intelligence. This ensures that security assessments are continuous rather than point-in-time, allowing for the rapid identification of vulnerabilities in both traditional code and complex ML architectures.

Domain 7: Security Operations

In the Security Operations Center (SOC), AI is a force multiplier. This domain focuses on the integration of AI and ML into Security Orchestration, Automation and Response (SOAR) platforms. The CISSP Exam Outline addresses how to manage "Alert Fatigue" by using AI to correlate disparate events and provide high-fidelity context to security analysts, allowing for faster incident response.

Operationally, we also cover the "Security of AI" during the production phase. This includes monitoring for "Model Drift"—where an AI's performance degrades over time—and responding to live adversarial attacks. By blending traditional incident response with AI-specific monitoring, CISSPs ensure that the organization's operational resilience keeps pace with the speed of automated threats.

Domain 8: Software Development Security

As AI transforms how code is written, Domain 8 has evolved to secure the modern development lifecycle. The CISSP Exam Outline incorporates the use of AI-assisted coding tools, focusing on the risks of "hallucinated" vulnerabilities or the accidental inclusion of insecure code snippets generated by LLMs. The focus is on integrating automated AI security testing into the CI/CD pipeline to catch these flaws before they reach production.

Additionally, this domain addresses the security of the software supply chain as it pertains to ML libraries and frameworks. Professionals are tasked with identifying and mitigating "Model Hijacking" or "Inference Attacks" that target the software layer. By embedding AI considerations into the Software Development Life Cycle (SDLC), we ensure that developers can leverage the efficiency of AI without compromising the integrity of the finished product.

Based on CCSP Exam Outline Effective August 1, 2026

As cloud environments become the primary infrastructure for hosting and training LLM and ML pipelines, the Certified Cloud Security Professional (CCSP) has evolved to meet these challenges. The CCSP Exam Outline explicitly integrates AI security across all six domains, ensuring that cloud architects and engineers can design, deploy and manage AI-driven services without compromising data sovereignty or architectural integrity.

Domain 1: Cloud Concepts, Architecture and Design

In the cloud architecture domain, AI integration begins with the evaluation of Cloud Service Provider (CSP) capabilities for hosting specialized AI workloads. Architects must now understand the shared responsibility model as it applies to AI-as-a-Service (AIaaS), specifically identifying where the provider's responsibility for model infrastructure ends and the customer's responsibility for model configuration and data begins. This includes designing for high-performance compute enclaves that can handle the massive throughput required for AI training while maintaining logical isolation.

Furthermore, this domain addresses the integration of AI into the cloud design process itself. The CCSP Exam Outline emphasizes the use of "Infrastructure as Code" (IaC) to deploy resilient AI environments and the application of cloud-native security design principles to mitigate the risks of model inversion and extraction. By building security into the foundation of the cloud-AI stack, CCSPs ensure that intelligent services are scalable, compliant and defensible from the moment they are provisioned.

Domain 2: Cloud Data Security

Data is the most critical asset in the cloud-AI ecosystem, and Domain 2 focuses on its protection throughout the AI lifecycle. Integration efforts focus on the security of massive "Data Lakes" and training sets, implementing advanced discovery and classification tools that use AI to identify sensitive information before it enters an ML pipeline. The CCSP Exam Outline addresses the unique challenges of data sovereignty and jurisdictional risk when AI training data is processed across multiple cloud regions.

Additionally, this domain incorporates technical controls such as homomorphic encryption and differential privacy to protect data during the inference phase. Practitioners manage the "Data Remanence" risks associated with ephemeral storage used by AI nodes, ensuring that sensitive training residuals are completely purged. This ensures that the use of cloud-based AI does not inadvertently lead to data leakage or a violation of global privacy regulations.

Domain 3: Cloud Platform and Infrastructure Security

This domain focuses on the hardened infrastructure required to run AI safely in the cloud. The CCSP Exam Outline integrates AI by addressing the security of the virtualization and containerization layers that host ML models. CCSPs are tasked with implementing specialized micro-segmentation to isolate AI training clusters and utilizing cloud-native hardware security modules (HSMs) to protect the cryptographic keys used for model signing and data encryption.

Moreover, we consider the role of AI in physical and logical infrastructure defense. This includes the administration of cloud-based DDoS protection and Web Application Firewalls (WAF) that leverage machine learning to detect and block sophisticated "Low and Slow" attacks targeting AI endpoints. By securing the underlying cloud platform, we ensure that the compute resources powering AI remain available and resilient against adversarial manipulation.

Domain 4: Cloud Application Security

As cloud applications increasingly leverage AI APIs, Domain 4 addresses the security of these integrations. The CCSP Exam Outline incorporates the secure development lifecycle (SDLC) for AI-driven cloud apps, focusing on the risks of insecure API calls and the potential for "Inference Attacks" at the application layer. Practitioners are able to implement robust input validation and rate limiting to defend against prompt injection and automated resource exhaustion targeting AI-based features.

This domain also covers the security of the software supply chain, specifically regarding the inclusion of third-party ML libraries and pre-trained models. CCSPs perform security testing on AI-augmented applications, ensuring that automated "hallucinations" or logic flaws do not introduce new vulnerabilities into the production environment. This ensures that cloud applications remain secure even as they become more autonomous and complex.

Domain 5: Cloud Security Operations

In the cloud SOC, AI serves as a vital tool for managing the sheer scale of cloud-generated logs. Integration within Domain 5 focuses on using AI and ML for advanced threat hunting and event correlation across multi-cloud environments. Candidates are aware that cloud-native SIEM/SOAR platforms can automate the response to common cloud threats, allowing human analysts to focus on high-complexity AI-driven attacks.

Operationally, this domain also addresses the "Continuous Monitoring" of AI performance to detect security-related "Model Drift." CCSPs are responsible for maintaining the operational baseline of AI services, ensuring that any deviation in model output is investigated as a potential security incident. This proactive approach ensures that cloud-based AI systems remain reliable and secure throughout their operational lifespan.

Domain 6: Legal, Risk and Compliance

The final domain addresses the complex regulatory landscape surrounding cloud-based AI. The CCSP Exam Outline integrates AI by focusing on the legal implications of automated data processing and the requirements for "Explainability" under frameworks like the GDPR or the EU AI Act. CCSPs understand how to conduct Cloud Data Life Cycle audits that specifically account for how AI models process and store data across international borders.

Risk management in this domain involves assessing the "Vendor Risk" of AI service providers, ensuring that their security controls and ethical guidelines align with organizational standards. We also cover the role of eDiscovery and digital forensics in the cloud when AI is involved, ensuring that practitioners can effectively investigate incidents and provide auditable evidence of compliance in a rapidly evolving legal environment.

Based on CGRC Exam Outline Effective June 15, 2024

The Certified in Governance, Risk and Compliance (CGRC) is the definitive certification for professionals who manage the intersection of security, privacy and organizational strategy. In an era where Artificial Intelligence (AI) drives critical business decisions, the CGRC Exam Outline has been updated to provide a robust framework for governing intelligent systems. By embedding AI-specific tasks and subtasks throughout the Risk Management Framework (RMF), the CGRC ensures that professionals can navigate the complexities of algorithmic transparency, "black box" risk and the rapidly evolving global regulatory landscape for AI.

Domain 1: Security and Privacy Governance, Risk Management and Compliance Program

Governance in the age of AI requires establishing dedicated oversight boards to manage algorithmic transparency and the ethical use of autonomous agents. Within this domain, we integrate AI by adapting traditional risk management principles to account for the unique, non-deterministic decision-making of machine learning models. This involves mapping complex, overlapping global requirements—such as the NIST AI Risk Management Framework (AI RMF) and ISO/IEC 42001—into existing corporate compliance tracking tools to ensure a unified and ethical approach to AI adoption.

Furthermore, this domain addresses the information lifecycle specifically for AI, focusing on the technical challenge of "machine unlearning" when sensitive data must be purged from trained models. By utilizing AI-driven analytics to define strict privacy guardrails for LLM training, CGRC professionals ensure that the organization's governance program protects intellectual property and data confidentiality while enabling responsible innovation.

Domain 2: Scope of the System

Traditional scoping methodologies have been expanded to capture the sprawling and continuous nature of modern machine learning data pipelines. In Domain 2, the CGRC Exam Outline integrates AI by requiring professionals to identify all embedded algorithms, including those hidden within commercial-off-the-shelf (COTS) software. This involves documenting the probabilistic nature of AI subsystems and utilizing AI-driven automated mapping tools to maintain dynamic system descriptions that keep pace with rapidly evolving cloud-native infrastructures.

Accurate scoping also requires a clear delineation between the foundational AI model training environment and the active inference endpoints. By using Natural Language Processing (NLP) tools to extract and verify scoping boundaries from technical documentation, CGRC practitioners ensure that the assessment boundary is precisely defined. This prevents "scope creep" and ensures that the security and privacy obligations of the AI system are clearly understood by all stakeholders.

Domain 3: Selection and Approval of Framework, Security and Privacy Controls

Selecting the right defenses for an AI system requires integrating specialized overlays, such as the CSA AI Controls Matrix, into traditional framework baselines. This domain emphasizes the use of AI to automate the dynamic mapping and selection of controls across complex, hybrid architectures. Practitioners understand how to identify inherited AI security controls from major Cloud Service Providers (CSPs), ensuring that foundational models like those in AWS Bedrock or Azure OpenAI are properly accounted for in the system's security plan.

Moreover, the integration focuses on tailoring controls to mitigate AI-specific threats, including prompt injection and adversarial data poisoning. By leveraging AI algorithms to recommend optimal control selections based on a model's unique risk profile, CGRC professionals can rapidly document complex control inheritance hierarchies. This ensures that the selected safeguards are not only compliant with international standards but are also technically effective against modern algorithmic attacks.

Domain 4: Implementation of Security and Privacy Controls

The implementation phase addresses the deployment of AI-native security controls seamlessly into distributed machine learning pipelines. This domain focuses on the use of intelligent Infrastructure as Code (IaC) to automate the provisioning and configuration of these controls, ensuring consistency across hyper-scaled environments. Practitioners are tasked with developing strategies that implement these safeguards without introducing unacceptable latency into real-time AI inference endpoints, balancing security with operational performance.

Compliance is further strengthened by aligning implementation with emerging international requirements, such as the EU AI Act. CGRC professionals ensure that privacy controls are technically robust enough to prevent unauthorized model retraining on protected tenant data. By utilizing predictive analytics to model implementation timelines and funding requirements, they ensure that the organization's AI infrastructure is deployed in a way that is both fiscally responsible and fundamentally secure.

Domain 5: Assessment/Audit of Security and Privacy Controls

Auditing in an AI-driven environment shifts from manual inspections to the use of AI-powered audit tools that can correlate compliance evidence across vast cloud landscapes. This domain integrates the assessment of "black-box" machine learning models for risks like algorithmic bias, data poisoning and hallucinations. Practitioners understand how to use predictive AI to accurately scope the boundaries of auto-scaling ML infrastructure, ensuring that the audit reflects the actual operational state of the system.

Accountability is a cornerstone of this domain, specifically regarding the autonomous decisions made by an AI system during an audit period. The exam outline aligns specific responsibilities to AI Ethics Officers and ML Engineers, ensuring that the audit process captures the technical and ethical dimensions of AI performance. By automating the scheduling and resource allocation for these complex audits, CGRC professionals can provide authorizing officials with high-fidelity data regarding the system's true compliance posture.

Domain 6: System Compliance

The authorization of an AI system involves navigating the inherent uncertainties of generative AI through formal risk acceptance criteria. In Domain 6, the CGRC Exam Outline integrates the use of AI governance tools to automate the generation and submission of massive compliance authorization packages, such as System Security Plans (SSPs). Utilizing NLP to cross-reference thousands of pages of privacy documentation, practitioners can identify inconsistencies and streamline the review process for authorizing officials.

Furthermore, this domain addresses the specialized documentation required for intelligent systems, including algorithmic bias audits and training data provenance logs. By using AI orchestration to route these documents through complex, multi-tier stakeholder approval workflows, CGRC professionals ensure that the final authorization is based on a transparent and comprehensive understanding of the system's risks. This ensures that even the most advanced AI deployments satisfy the rigorous standards for organizational risk acceptance

Domain 7: Compliance Maintenance

Maintaining compliance for an AI system requires a transition to AI-driven Continuous Control Monitoring (CCM) that can handle the rapid change lifecycle of MLOps. This domain integrates the governance of continuous deployments, treating the update of ML weights as formal system changes. Practitioners understand how to use AI to predict the "blast radius" and compliance impact of a proposed architectural change before it is approved, preventing unauthorized drift from the authorized security baseline.

Finally, the CGRC Exam Outline addresses the role of AI in autonomously assessing how updates to foundational models might alter the system's compliance posture or introduce new biases. By evaluating how shifts in training data demographics impact regulatory requirements, CGRC professionals ensure that the system remains compliant throughout its operational life. This proactive approach to maintenance ensures that the organization can leverage the latest AI innovations without compromising its long-term security and compliance goals.

Based on CSSLP Exam Outline Effective September15, 2023

The Certified Secure Software Lifecycle Professional (CSSLP) represents the gold standard in application security excellence. As AI transforms the software development landscape, the CSSLP Exam Outline has been modernized to address the security of the entire lifecycle—from initial concept to supply chain management—specifically as it pertains to AI-driven and AI-integrated applications. This ensures that software professionals can design, build and maintain applications that leverage machine learning safely, defending against the unique, non-deterministic risks that AI brings to the software stack.

Domain 1: Secure Software Concepts

In the foundational domain of the CSSLP, AI integration focuses on how generative AI and LLMs fundamentally alter traditional software security boundaries. Professionals must now redefine core pillars of information security concepts to address vulnerabilities unique to machine learning, such as data poisoning and model inversion. This involves applying secure design principles to probabilistic outputs, ensuring that the software's foundational concepts are resilient to the non-deterministic behavior of integrated AI components.

Furthermore, this domain addresses the risk of AI-driven systems leaking sensitive proprietary code or Personally Identifiable Information (PII) through inadvertent memorization during the training phase. By utilizing AI for automated key management and continuous enforcement of security primitives, CSSLPs ensure that security is not just an overlay but is embedded into the very logic of the software's core concepts from the beginning.

Domain 2: Secure Software Lifecycle Management

As development teams adopt machine learning, the CSSLP focuses on the transition from traditional DevSecOps to MLSecOps. This domain integrates the management of non-linear, continuous retraining loops of AI models into standard Agile sprints. Security managers oversee the secure use of generative AI coding assistants, ensuring that these tools accelerate the lifecycle without introducing "hallucinated" vulnerabilities or bypassing critical security gates.

Operationally, this domain emphasizes the adoption of emerging frameworks like the OWASP Top 10 for LLMs alongside traditional application security standards. Professionals are tasked with establishing explicit build-break criteria for CI/CD pipelines when an AI model's bias or hallucination rate exceeds safety thresholds. This ensures that the entire lifecycle—from strategic roadmaps for post-quantum cryptography to automated security ticketing—is capable of managing the speed and complexity of AI-driven development.

Domain 3: Secure Software Requirements

Requirement engineering now includes defining strict boundaries for the integration of third-party LLMs and AI microservices. Professionals must establish functional requirements to prevent AI models from executing unauthorized autonomous actions and define non-functional requirements for acceptable hallucination rates and algorithmic bias thresholds. This integration ensures that the business use cases for AI are clearly scoped, preventing "function creep" that could lead to unmanaged security risks.

Additionally, this domain leverages AI to enhance the requirements process itself, utilizing Natural Language Processing (NLP) to autonomously parse and validate complex requirement documents. By identifying "abuse cases" specifically targeting embedded chatbots or inference engines, CSSLPs ensure that security requirements are robust enough to withstand adversarial attempts to jailbreak or manipulate the application's intelligence layer.

Domain 4: Secure Software Architecture and Design

In the architecture domain, the focus is on designing resilient systems that decouple core application logic from unpredictable AI inference engines. Architects understand how to design Zero Trust boundaries around highly sensitive vector databases and AI processing clusters, ensuring that the integration of machine learning does not compromise the broader system's security architecture. This includes utilizing AI-driven threat modeling to autonomously map evolving attack surfaces during the design phase.

Architects also adapt high-level frameworks like SABSA to account for the probabilistic nature and unique supply chain of machine learning. By securing distributed computing nodes that utilize federated learning and managing the high-volume message queuing required for AI inference, CSSLPs ensure that the software architecture is both performant and defensible against the specific failure modes of AI systems.

Domain 5: Secure Software Implementation

Implementation now centers on the secure use of AI-assisted coding and the defense of embedded machine learning algorithms. Professionals must adhere to secure coding standards explicitly designed to mitigate prompt injection and ensure the provenance of AI-generated code snippets. This domain focuses on the rigorous sanitization of all inputs sent to LLMs, utilizing NLP-driven validation to semantically understand and block malicious payloads that traditional input filters might miss.

Furthermore, the implementation phase addresses the use of AI to autonomously translate declarative security policies into enforced imperative code at runtime. By evaluating how these declarative boundaries can restrict the autonomous actions of AI agents, CSSLPs ensure that the code implementation is robust enough to handle high-throughput AI workloads while preventing the introduction of vulnerabilities during the transition from design to production.

Domain 6: Secure Software Testing

Software testing has evolved from purely deterministic methods to the probabilistic testing required for AI model outputs. Within this domain, professionals integrate specific testing phases for embedded ML models to evaluate them for bias, drift, toxicity and susceptibility to adversarial evasion. This includes utilizing AI to autonomously generate complex, edge-case security test scripts at machine speed, ensuring that the testing coverage keeps pace with the complexity of the application.

Testing also encompasses the functional logic boundaries of autonomous AI agents, validating that they cannot execute critical commands outside of their intended scope. By stress-testing AI models against token limits and evaluating them against the NIST AI Risk Management Framework (AI RMF), CSSLPs ensure that the application fails securely and remains reliable even when the AI component experiences performance degradation or adversarial interference.

Domain 7: Secure Software Deployment, Operations and Maintenance

This domain focuses on the unique operational risks of deploying non-deterministic AI models into deterministic software environments. We integrate "AIOps" to autonomously monitor and maintain software infrastructure, allowing for self-healing systems that can respond to operational anomalies in real-time. Maintenance now includes training administrators to identify signs of algorithmic drift and adversarial AI attacks occurring in production environments.

Practitioners are also tasked with creating isolated, sandboxed staging environments designed specifically to test the behavior of updated AI model weights before full deployment. By managing the legal and copyright risks of AI-generated code and ensuring compliance with emerging laws like the EU AI Act, CSSLPs ensure that the ongoing operation and maintenance of AI-driven software remains transparent, compliant and secure.

Domain 8: Secure Software Supply Chain

The software supply chain has expanded to include the unique risks of the AI ecosystem, such as reliance on external foundational models and massive public datasets. This domain focuses on the evolution of traditional SBOMs into "AI-BOMs" (AI Bill of Materials), which track model weights, training datasets and ML libraries. Professionals understand how to establish rigorous security criteria for selecting open-source AI models, preventing the integration of malicious backdoors or biased components.

Finally, this domain utilizes AI tools to autonomously map and secure deeply nested, complex supply chains. By quantifying the risk of "model poisoning" and assessing the inherent bias of pre-trained open-source models, CSSLPs protect the integrity of the software from upstream vulnerabilities. This holistic approach ensures that every component—from third-party libraries to the training data itself—satisfies the organization's security and ethical standards.

Based on **ISSAP Exam Outline** Effective August 1, 2025

The Information Systems Security Architecture Professional (ISSAP) certification represents the pinnacle of security design. As organizations transition to AI-native infrastructures, the ISSAP Exam Outline has evolved to ensure that architects can design complex, resilient environments that treat AI as both a powerful defensive asset and a high-value protected surface. By embedding AI considerations into the architecture lifecycle, the ISSAP Exam Outline ensures that senior architects can align business-driven AI innovation with the most rigorous security engineering standards.

Domain 1: Governance, Risk and Compliance (GRC)

In the realm of identity architecture, the ISSAP Exam Outline addresses the complex challenge of managing identities for autonomous AI agents and automated service accounts. Architects are tasked with designing "Identity-as-a-Service" (IDaaS) frameworks that enforce the Principle of Least Privilege for non-human entities, ensuring that AI systems cannot escalate their own permissions as they navigate the enterprise. This domain emphasizes the architectural integration of behavioral biometrics and AI-driven adaptive authentication to move organizations toward a true Zero Trust posture.

Furthermore, we incorporate the use of AI to enhance the orchestration of access controls across hybrid and multi-cloud environments. Architects understand how to design automated provisioning and de-provisioning workflows that react in real-time to detected anomalies in user behavior. By placing AI at the heart of identity orchestration, the ISSAP ensures that the architecture can scale to meet the demands of an increasingly automated and decentralized workforce.

Domain 2: Security Architecture Modeling

The security architecture modeling domain focuses on the architectural design of the "Intelligent SOC." We integrate AI by addressing the infrastructure requirements for Security Orchestration, Automation and Response (SOAR) platforms and AI-driven Security Information and Event Management (SIEM) systems. Architects understand how to design high-throughput data pipelines that can feed massive volumes of telemetry into ML-powered correlation engines without introducing latency or data loss.

Beyond defensive automation, this domain covers the architecture required to monitor and defend AI models themselves. This includes the placement of specialized "AI Firewalls" and monitoring probes designed to detect prompt injection or model evasion attempts. By architecting a unified visibility layer, the ISSAP ensures that security operations can maintain a cohesive defense strategy across both traditional assets and sophisticated machine learning pipelines.

Domain 3: Infrastructure and System Security

Infrastructure architecture accounts for the specialized, high-performance compute environments required for AI training and inference. The integration focuses on the use of "Hardware-Rooted Trust" and Trusted Execution Environments (TEE) to protect sensitive model weights and training data at the physical layer. Architects are tasked with designing secure enclaves and micro-segmentation strategies that isolate AI workloads, preventing an adversary from leveraging a compromised AI interface to move laterally through the data center.

We also address the role of AI in defending the network and physical perimeter. Architects understand how to integrate AI-powered network sensors and Software-Defined Perimeter (SDP) solutions that can dynamically adjust to emerging threats. By building a resilient, AI-aware infrastructure, the ISSAP ensures that the foundational layers of the enterprise are capable of supporting the massive computational and security demands of modern intelligence systems.

Domain 4: Identity and Access Management (IAM) Architecture

At the architectural level, GRC integration involves designing systems that are "secure and compliant by design." This domain focuses on the architectural implementation of the NIST AI Risk Management Framework (AI RMF) and global regulatory requirements like the "Right to Explanation." Architects understand how to design transparent system architectures that provide auditable logs of AI decision-making, ensuring that automated processes can be verified by human overseers and legal auditors.

Additionally, this domain addresses the risk architecture for third-party and supply-chain AI integrations. We emphasize the design of "Vendor Risk Architecture" that evaluates the security posture of AI service providers before they are integrated into the enterprise. By embedding risk management directly into the architectural blueprints, the ISSAP ensures that AI adoption remains within the organization's established risk appetite and legal boundaries.

Based on ISSEP Exam Outline Effective August 1, 2025

The Information Systems Security Engineering Professional (ISSEP) certification is where rigorous systems engineering meets advanced cybersecurity. As AI becomes a primary component of mission-critical systems, the ISSEP Exam Outline ensures that security engineers can mathematically and architecturally validate the integrity of AI-driven components. By embedding AI into the systems engineering lifecycle, the ISSEP Exam Outline ensures that the "intelligence" of a system is as defensible and predictable as its hardware and software.

Domain 1: Systems Security Engineering Foundations

In the foundational domain of the ISSEP, AI integration focuses on the mathematical and theoretical validation of machine learning models within a system's security architecture. Security engineers are now tasked with applying formal methods to AI components, ensuring that "Algorithmic Integrity" is maintained throughout the system lifecycle. This involves integrating AI-specific threat modeling—such as identifying vulnerabilities to adversarial evasion—directly into the initial design requirements to ensure that the system's baseline is resilient against non-deterministic threats.

Furthermore, this domain addresses the incorporation of AI into the "Security Design Principles." Engineers must now account for the unique computational and data-handling requirements of neural networks, ensuring that the integration of an AI sub-system does not violate established security enclaves or trust boundaries. By treating the ML model as a high-value engineering asset, ISSEPs ensure that the foundation of the system is built to withstand both traditional and AI-augmented attack vectors.

Domain 2: Risk Management

Risk management for the security engineer has evolved to include the probabilistic risks inherent in AI. Within this domain, the ISSEP Exam Outline integrates methodologies for assessing "Model Robustness" and the potential for logic drift in production environments. Practitioners utilize AI-driven assessment tools to perform continuous, automated vulnerability scanning of complex system interdependencies, shifting from point-in-time audits to a real-time understanding of the system's risk posture.

Moreover, the ISSEP Exam Outline focuses on the engineering of "Explainable AI" as a core risk-mitigation strategy. By designing systems that provide transparent, auditable pathways for their decisions, engineers can provide the necessary evidence for security authorizations and certifications. This integration ensures that even the most complex deep-learning components can meet the rigorous documentation and verification standards required for high-assurance environments.

Domain 3: Security Planning and Engineering

This domain focuses on the "Build" phase, where AI is integrated into the secure systems development lifecycle (SSDLC). Security engineers are tasked with designing secure data ingestion pipelines that prevent "Data Poisoning" during the training phase. The integration emphasizes the use of hardware-rooted trust, such as Trusted Execution Environments (TEE), to protect AI model weights and inference logic from unauthorized access or tampering at the physical layer.

Additionally, the ISSEP Exam Outline addresses the secure integration of AI APIs and third-party ML libraries. Engineers must evaluate the "provenance" of pre-trained models, ensuring that the software supply chain for the system's intelligence is as secure as the code itself. By embedding AI security requirements into the functional design specifications, ISSEPs ensure that the finished system is not only intelligent but also architecturally sound and defensible.

Domain 4: Systems Security Implementation, Verification and Validation

Verification and Validation are critical in engineering, and this domain includes the testing of AI outputs against established security policies. Integration involves developing "Adversarial Testing" protocols where engineers attempt to trick or bypass AI-driven controls to verify their resilience. This ensures that the system's automated responses are consistent, predictable and do not introduce "hallucinations" that could compromise mission success.

By utilizing machine learning to parse massive amounts of testing data and system logs, engineers can more effectively identify edge cases and logic flaws that traditional testing might miss. This dual-sided integration ensures that as systems become more complex, the engineering tools used to validate them remain equally sophisticated and effective.

Domain 5: Secure Operations, Change Management and Disposal

The final domain addresses the long-term sustainability of AI-integrated systems. Engineers are tasked with designing "Continuous Monitoring" frameworks that specifically track model performance and integrity over time. This includes establishing automated triggers for model retraining or system rollback if "Concept Drift" or a security compromise is detected. This ensures that the system remains within its authorized security parameters throughout its entire operational lifespan.

Maintenance also involves the secure patching and updating of AI models in the field. ISSEPs must design secure delivery mechanisms for large-scale model updates, ensuring that the integrity of the "intelligence" is verified at every step of the update process. By integrating AI into the maintenance cycle, the ISSEP Exam Outline ensures that the system remains resilient against evolving threats and that the "human-in-the-loop" remains empowered to oversee and control autonomous system functions.

Based on **ISSMP Exam Outline** Effective August 1, 2025

The Information Systems Security Management Professional (ISSMP) certification is designed for leaders who bridge the gap between technical security and organizational strategy. As AI becomes a core driver of business innovation, the ISSMP Exam Outline includes critical management competencies for overseeing AI and ML. From establishing ethical governance to architecting resilient, AI-powered security operations, the integrated ISSMP Exam Outline ensures that senior security managers can lead their organizations through the complexities of the algorithmic era while maintaining a robust and compliant security posture.

Domain 1: Leadership and Organizational Management

In the realm of strategic leadership, AI integration focuses on guiding executive teams through the secure adoption of machine learning into core business functions. This domain emphasizes the establishment of ethical AI governance models that ensure algorithmic transparency and accountability for automated decision-making. Security managers are tasked with fostering a culture that encourages secure AI experimentation while preventing the risks associated with "Shadow AI"—the unauthorized use of public AI tools that could lead to data leakage or reputational damage.

Furthermore, the domain addresses the alignment of security vision with the unique demands of the AI era. This involves updating organizational missions to protect intellectual property in the age of generative AI and embedding security checkpoints directly into data science workflows. By redefining organizational processes to account for the speed of AI-assisted operations, ISSMPs ensure that security remains an enabler of innovation rather than a bottleneck to progress.

Domain 2: Systems Lifecycle Management

The transition from traditional, deterministic systems to continuous, probabilistic machine learning pipelines is a primary focus of this domain. Security managers understand how to oversee "MLSecOps" lifecycles, integrating automated, AI-driven security testing and dynamic validation at every phase. The integration introduces explicit decision gates within the development process to halt deployments if an AI model exceeds established thresholds for bias or "hallucinations," ensuring that only safe and reliable models reach production.

Additionally, this domain includes the use of AI to enhance the management of enterprise architectures. Practitioners are expected to leverage autonomous discovery and classification tools to categorize massive volumes of unstructured data that feed into AI systems. By managing the continuous feedback loops required to secure and retrain active models, ISSMPs ensure that the security of an application is maintained even as the underlying machine learning logic evolves over time.

Domain 3: Risk Management

Risk management has evolved from a static, point-in-time assessment to a dynamic, real-time discipline. The ISSMP Exam Outline integrates globally recognized frameworks, such as NIST AI Risk Management Framework (AI RMF) and ISO/IEC 42001 alongside traditional strategies, enabling managers to govern the ethical use and procurement of generative AI. By utilizing predictive machine learning models for risk scoring, organizations can transition to a more proactive posture, identifying and mitigating threats to proprietary model weights and algorithms before they are exploited.

The scope of risk management also expands to include the complex ecosystem of third-party LLMs and the massive data lakes required for continuous learning. ISSMPs are to engage new stakeholders—such as Chief Data Officers and AI Ethics Committees—to define acceptable error thresholds and mitigate the reputational impact of AI-driven decisions. This ensures that the organizational risk strategy is comprehensive enough to cover both human and algorithmic vulnerabilities.

Domain 4: Security Operations

In the modern Security Operations Center (SOC), AI is both a powerful defensive tool and a new attack surface. This domain focuses on managing "AIOps" to autonomously detect and remediate threats, moving beyond reactive alert triage to proactive, AI-driven threat hunting. Security managers are responsible for overseeing the operational response to novel adversarial attacks, such as prompt injection and data poisoning, which specifically target the logic of the organization's AI models.

Operational integration also involves the use of Generative AI to rapidly author and update complex playbooks and runbooks. Data Scientists and ML Engineers are key stakeholders in the incident response process. By structuring teams to include AI Security Analysts who specialize in reverse-engineering algorithmic evasion tactics, ISSMPs ensure that the SOC remains resilient in the face of machine-speed attacks.

Domain 5: Contingency Management

Resiliency planning must now account for the massive scale and specialized infrastructure required by modern AI. This domain integrates AI by using generative modeling to simulate and draft highly complex disaster recovery (DR) plans. Security managers understand how to architect contingency strategies that prioritize the restoration of critical data ingestion pipelines and massive vector databases, ensuring that the predictive accuracy of business-critical AI does not degrade during a disruption.

The ISSMP Exam Outline also addresses the unique challenge of "Model Drift" as a continuity risk. ISSMPs must analyze the resiliency of cloud architectures hosting massive LLMs and establish Recovery Time Objectives (RTO) that account for the extreme computational resources needed to retrain a model from scratch. By utilizing predictive AI to model the business impact of various disaster scenarios, practitioners can ensure that organizational resiliency keeps pace with the demands of an AI-dependent enterprise.

Domain 6: Law, Ethics and Security Compliance Management

The final domain navigates the rapidly shifting legal landscape, including global regulations like the EU AI Act and emerging standards for algorithmic liability. Integration efforts focus on managing the conflict between data minimization requirements and the massive data ingestion needs of continuous machine learning. ISSMPs understand how to implement dynamic data routing to comply with localized trans-border data flow restrictions when training centralized enterprise LLMs.

Ethics and compliance also cover the "Right to Explanation," where users must be informed of the logic behind automated profiling. Subtasks within this domain address the intellectual property risks of using open-source AI models and the necessity of maintaining conversational prompt histories in compliance with privacy laws. By ensuring that AI systems are transparent, unbiased and respectful of user privacy, the ISSMP serves as the guardian of the organization's legal and ethical integrity in an automated world.

Additional Examination Information

Review Exam Outlines and Experience Requirements

All exam candidates are advised to review the exam outline and experience requirements thoroughly before selecting and pursuing their exam. This document is not intended to replace the ISC2 certification exam outline as the source for guidance on the domains and topics candidates can expect to encounter during their certification exam. Exam outlines can be found on our certification web pages at **www.ISC2.org/certifications**.

Examination Policies and Procedures

ISC2 recommends that candidates review exam policies and procedures prior to registering for the examination. Read the comprehensive breakdown of this important information at **ISC2.org/register-for-exam**.

About ISC2

ISC2 is the world's leading member organization for cybersecurity professionals, driven by our vision of a safe and secure cyber world. Our more than 265,000 certified members, and associates, are a force for good, safeguarding the way we live. Our award-winning certifications – including cybersecurity's premier certification, the CISSP® – enable professionals to demonstrate their knowledge, skills and abilities at every stage of their careers. ISC2 strengthens the influence, diversity and vitality of the cybersecurity profession through advocacy, expertise and workforce empowerment that accelerates cyber safety and security in an interconnected world. Our charitable foundation, the **Center for Cyber Safety and Education**, helps create more access to cyber careers and educates those most vulnerable. Learn more, get involved or become an ISC2 Candidate to build your cyber career at **ISC2.org**. Connect with us on **X**, **Facebook** and **LinkedIn**.