GSDC
Global Skill Development Council

GENERATIVE AI CYBERSECURITY

Interview Preparation Guide

www.gsdcouncil.org

Q1. How would you explain Generative AI to a non-technical cybersecurity stakeholder?

Generative AI is artificial intelligence that doesn't just analyze existing data — it creates new content by learning patterns from massive datasets. In cybersecurity, this works in two directions simultaneously. Defenders use it to generate threat reports, analyze malware code, and automate incident responses faster than any human team. Attackers use the same technology to write convincing phishing emails, generate polymorphic malware, and create deepfake impersonations. The key message for non-technical stakeholders is this: Gen AI doesn't change the fundamentals of cybersecurity — it dramatically accelerates everything on both sides of the battle.

Q2. What is the difference between traditional ML and Generative AI in a security context?

Traditional machine learning in cybersecurity is primarily discriminative — it classifies, detects, and predicts based on labeled training data. It answers "is this malicious or not?" Generative AI goes further — it creates new content, explains its reasoning in plain language, synthesizes intelligence from unstructured sources, and adapts to novel situations it has never explicitly seen.

Traditional ML

Catches known malware variants through pattern matching — signature-based, discriminative, labeled data dependent.

Generative AI

Analyzes unknown malware code, explains what it does, and predicts how it might evolve — all in natural language an analyst can act on immediately.

Q3. What is a Large Language Model and why does it matter for cybersecurity professionals?

A Large Language Model is an AI system trained on enormous volumes of text data that can understand and generate human language at a level that enables genuine reasoning tasks. For cybersecurity professionals, LLMs matter because they power the new generation of security tools — Microsoft Security Copilot, CrowdStrike Charlotte AI, SentinelOne Purple AI — that allow analysts to investigate incidents through conversation rather than complex query languages.

Instead of writing a 20-line SIEM query, an analyst asks "show me all lateral movement activity from the finance segment in the last 48 hours" and gets the answer.

LLMs also power automated threat report generation, vulnerability explanation, and policy drafting — removing significant administrative burden from security teams.

Q4. What is RAG and how is it used in cybersecurity tools?

RAG stands for Retrieval-Augmented Generation. It is an AI architecture that retrieves relevant information from a specific knowledge base before generating a response — rather than relying solely on what the model learned during training. In cybersecurity, RAG is critical because it allows AI tools to answer questions grounded in your organization's actual data: your specific CVE inventory, your past incident history, your internal threat intelligence, your security policies.

Without RAG: A security AI assistant gives generic answers based on public training data — a generic security chatbot.

With RAG: It gives specific answers based on your environment — a genuinely useful internal security assistant.

The difference between a generic security chatbot and a genuinely useful internal security assistant is almost always RAG.

Q5. What is prompt injection and why should cybersecurity teams care about it?

Prompt injection is an attack where malicious content embedded in data processed by an AI system manipulates the AI into ignoring its original instructions or revealing sensitive information. It is the AI equivalent of SQL injection — instead of injecting malicious database commands, attackers inject malicious instructions into inputs the AI processes.

Risk #1: Vulnerable AI Deployments

Any AI tool your organization deploys that processes external content — emails, web pages, user inputs — is potentially vulnerable to prompt injection attacks that could cause the AI to behave maliciously.

Risk #2: Targeted Payloads

Attackers are increasingly embedding prompt injection payloads in documents and emails specifically targeting AI-assisted security tools, attempting to manipulate the AI analyst into missing or misreporting threats.

Q6. What is AI hallucination and what specific risks does it create in a security context?

Hallucination occurs when an AI model generates confident, plausible-sounding but factually incorrect information. In general contexts, hallucinations are inconvenient. In cybersecurity, they can be dangerous. Specific risks include:

- An AI generating an incorrect CVE description that leads a team to misunderstand the severity of a vulnerability
- An AI incident report attributing an attack to the wrong threat actor, misdirecting the investigation
- An AI threat assessment confidently clearing a genuine threat as benign
- An AI suggesting a remediation step that doesn't actually address the real vulnerability

The mitigation is always the same — establish mandatory human verification checkpoints for all AI-generated security outputs before acting on them, especially for high-consequence decisions.

Q7. Explain the difference between an Indicator of Compromise (IoC) and an Indicator of Attack (IoA).

Indicator of Compromise (IoC)

Forensic evidence that a breach has already occurred — artifacts left behind by an attacker: malicious IP addresses, file hashes, domain names, registry keys modified by malware. IoCs are retrospective.

Indicator of Attack (IoA)

Behavioral — signals that an attack is in progress right now, based on what an attacker is doing rather than what they have left behind. Examples: a process spawning an unusual child process, PowerShell executing encoded commands, or a user account accessing resources it has never accessed before.

Gen AI is particularly valuable for IoA detection because behavioral patterns require contextual reasoning across many data points simultaneously — exactly what AI excels at and what rules-based systems consistently miss.

Q8. What is the MITRE ATT&CK framework and how does Gen AI interact with it?

MITRE ATT&CK is a globally recognized knowledge base of adversary tactics and techniques based on real-world attack observations. It is organized into 14 tactic categories — from Reconnaissance through Impact — each containing specific techniques and sub-techniques that attackers use. Gen AI interacts with ATT&CK in several important ways:

1. Automatic Mapping: Automatically maps detected security events to ATT&CK techniques, giving analysts a standardized way to understand and communicate threats.
2. Hunt Hypothesis Generation: Generates hunt hypotheses based on ATT&CK techniques relevant to specific threat actor groups.
3. Automated Reporting: Produces incident reports that include ATT&CK mapping automatically.
4. Natural Language Querying: Allows analysts to query security data using ATT&CK language — "show me all T1059 command and scripting interpreter activity" — without needing to know the underlying technical query syntax.

Q9. How does Gen AI improve threat intelligence operations compared to traditional approaches?

Traditional threat intelligence was bottlenecked by analyst capacity — the volume of intelligence sources, dark web forums, CVE databases, vendor reports, and OSINT feeds vastly exceeded what any team could meaningfully process. Gen AI removes that bottleneck in three ways:

Continuous Ingestion. Continuously ingests and synthesizes intelligence from hundreds of sources simultaneously, producing structured, readable intelligence reports without analyst involvement in the collection phase.

IoC Extraction. Extracts Indicators of Compromise from unstructured sources — blog posts, forum discussions, malware analysis reports — automatically.

ATT&CK Mapping. Maps threat actor behavior to MITRE ATT&CK techniques without manual analysis, giving teams a consistent, standardized picture of the threat landscape.

The result is that threat intelligence becomes continuous and comprehensive rather than periodic and selective.

Q10. What is predictive threat intelligence and how does AI enable it?

Predictive threat intelligence uses AI to forecast likely future attacks based on current threat actor behavior, historical attack patterns, and your organization's specific exposure profile. Rather than only telling you what threats exist today, AI analyzes patterns to tell you which attack vectors are most likely to be used against your specific industry, geography, and technology stack in the near term. It works by correlating:

- Threat actor campaign history
- Vulnerability exploitation timelines
- Geopolitical triggers
- Seasonal attack patterns

This is a level of multi-variable correlation impossible for human analysts at scale. The value is shifting security posture from reactive — responding after detection — to proactive — strengthening defenses against the attacks most likely to come before they arrive.

Q11. Walk me through how AI-powered anomaly detection works in a SOC environment.

- Baseline Learning: Establish normal behavior for users, devices, apps.
- Continuous Monitoring: Compare all activity against baselines in real time.
- Risk Scoring: Compute deviation scores with plain-English explanations.
- Analyst Alert: Flag anomalies with context: access, volume, times.

The critical difference from rule-based detection is that AI doesn't require a predefined rule for every attack scenario — it detects anything that deviates meaningfully from established normal behavior, including novel attacks no one has written a rule for yet.

Q12. What is User and Entity Behavior Analytics (UEBA) and why is AI essential for it?

UEBA is a security capability that profiles the behavior of users and entities — devices, applications, service accounts — and detects anomalies that suggest compromise, insider threat, or account takeover. AI is essential for UEBA because meaningful behavior profiling requires processing enormous volumes of data across many dimensions simultaneously.

Rule-Based System: Could detect a user logging in from a foreign country — a single-dimension check.

AI-Powered UEBA: Detects that the same user is logging in from an unusual location, accessing resources outside their normal pattern, at an unusual time, with an unusual application — and contextualizes this against a current threat intelligence feed showing active credential stuffing campaigns targeting your industry. That level of multi-dimensional contextual reasoning is only possible with AI.

Q13. How does Gen AI assist with threat hunting specifically?

Gen AI transforms threat hunting in four practical ways:

- Hunt Hypothesis Generation: Based on current threat intelligence and your environment's specific exposure, it suggests what to hunt for rather than requiring hunters to generate all hypotheses manually.
- Natural Language Querying: Hunters ask questions in plain English and AI translates them into the correct query language for the SIEM, EDR, or NDR platform, removing the technical barrier for less experienced hunters.
- Automated Hunt Playbooks: Executes automated hunt playbooks continuously in the background, running predefined hunts around the clock without analyst initiation.
- Long-Period Pattern Identification: Connects a suspicious event from months ago with a current alert that a human analyst would never link manually, surfacing the attacker's full footprint retrospectively.

Q14. What is the Cyber Kill Chain and how does AI accelerate defender response across its stages?

The Cyber Kill Chain is Lockheed Martin's model of a cyberattack across seven stages: Reconnaissance, Weaponization, Delivery, Exploitation, Installation, Command and Control, and Actions on Objectives. AI accelerates defender response at every stage:

- Reconnaissance: AI attack surface monitoring detects when attackers are probing your environment.
- Delivery: AI email analysis detects phishing attempts before they reach end users.
- Exploitation: AI-powered vulnerability management has already prioritized and flagged the vulnerabilities being exploited.
- Installation: AI behavioral detection identifies malware installation behavior without signatures.
- Command & Control: AI network anomaly detection flags unusual outbound communication.

The key insight for interviews: AI compresses the time between each kill chain stage from the defender's perspective — earlier detection, faster response, smaller blast radius.

Q15. How does Gen AI change malware analysis and what are its limitations?

Gen AI changes malware analysis by providing instant code explanation — an analyst uploads a malware sample or decompiled binary and receives a plain-English description of what the code does, what system calls it makes, what persistence mechanisms it uses, and what its likely objective is. This compresses analysis time from hours to minutes for standard samples. AI also automatically generates YARA detection rules from samples, classifies malware into known families, and detects obfuscation techniques.

The limitations are important:

- Obfuscation Evasion: AI can be fooled by sophisticated obfuscation specifically designed to confuse LLMs.
- Hallucination Risk: It may hallucinate incorrect technical details about what malware does.
- Novel Malware Families: It is less reliable for genuinely novel, previously unseen malware families with no comparable training examples.
- Human expert validation remains essential for high-confidence malware analysis, particularly for targeted, sophisticated samples.

Q16. What is polymorphic malware and why does it create challenges for traditional security tools?

Polymorphic malware changes its code signature with every infection while maintaining its malicious functionality. Traditional antivirus and detection tools work primarily through signature matching — comparing file hashes and code patterns against a database of known malicious signatures. Polymorphic malware defeats this approach because each variant has a unique signature that doesn't match anything in the database.

Traditional Approach (Fails): "Does this file match a known malicious signature?" — Polymorphic malware always produces a unique signature, defeating this check entirely.

Gen AI Approach (Succeeds): "Does this code behave in ways consistent with malicious intent, regardless of its specific signature?" — Processes spawned, registry keys modified, network connections made, and files accessed reveal malicious intent even when the code is completely novel.

Q17. How does AI improve vulnerability prioritization and why does this matter?

Most organizations have thousands of open vulnerabilities at any point in time — far more than can be patched simultaneously. Traditional prioritization relied primarily on CVSS score, but a high CVSS score doesn't mean a vulnerability is being actively exploited or that it is exploitable in your specific environment. AI improves prioritization by scoring vulnerabilities across multiple dimensions simultaneously:

- CVSS Score: Baseline severity rating as a starting point, not the sole criterion.
- Active Exploitation: Whether the vulnerability is being actively exploited in the wild right now.
- Public Exploit Availability: Whether a working exploit is publicly available to lower-skilled attackers.
- Asset Criticality: The criticality of the affected asset in your specific environment and its exposure.

Q18. What is adversarial machine learning and why is it a concern for security AI deployments?

Adversarial machine learning refers to attacks specifically designed to fool AI models into making incorrect predictions or classifications. In cybersecurity, this means crafting malicious inputs that cause AI security systems to misclassify threats as benign — allowing attacks to pass through AI-powered defenses undetected. Examples include:

- Crafting network traffic that statistically resembles legitimate traffic while performing a genuine attack
- Modifying malware code in ways that cause AI behavioral analysis to classify it as safe
- Poisoning the training data of AI security models so they learn to classify a specific attacker's techniques as normal

This is a critical concern because as organizations increasingly rely on AI-powered security tools, sophisticated attackers are specifically developing techniques to evade them — requiring security AI systems to be regularly tested for adversarial robustness.

Q19. What is SOAR and how does Gen AI make it significantly more capable?

SOAR stands for Security Orchestration, Automation, and Response — platforms that connect security tools and automate response workflows through predefined playbooks. Traditional SOAR is powerful but brittle: playbooks follow fixed decision trees that break when they encounter situations not anticipated by the designer. Gen AI makes SOAR dramatically more capable in three ways:

- Adaptive Playbook Execution: AI reads the context of each specific incident and chooses different response paths rather than following a fixed script.
- Natural Language Playbook Creation: Security engineers describe what they want to happen in plain English and AI generates the technical playbook, removing the coding barrier for automation.
- Intelligent Orchestration: AI decides in real time which tools to invoke, in which order, based on the evolving incident context — going far beyond the predetermined sequences that traditional SOAR supports.

Q20. Explain how you would use Microsoft Security Copilot in a real SOC investigation.

Microsoft Security Copilot is an AI security analyst assistant that integrates across the Microsoft security stack — Sentinel, Defender, Intune, Entra ID. In a real SOC investigation, I would use it as follows:

- "Summarize this incident and tell me what has happened so far." It pulls correlated data from Sentinel and gives an instant incident summary.
- "What is the most likely attack path based on this activity?" It maps the behavior to MITRE ATT&CK techniques and suggests the attacker's likely next steps.
- "Run a KQL query to find all other devices this user has authenticated to in the last 24 hours." It writes and executes the query without me knowing KQL syntax.

Throughout, I'm investigating in minutes what would traditionally take hours of manual correlation. The critical point: Copilot amplifies analyst capability — it doesn't replace the analyst's judgment about what the findings mean and what action to take.

Q21. What are the key SOC metrics that AI improves and how does it improve them?

The four most important SOC metrics that AI directly improves are:

- Mean Time to Detect (MTTD): AI continuously monitors all telemetry for behavioral anomalies — catching threats that would take human analysts days to identify through manual log review.
- Mean Time to Respond (MTTR): AI executes the first steps of incident response automatically — isolating endpoints, blocking IPs, disabling accounts — before a human analyst has even begun investigating.
- Alert Fatigue: AI triages, correlates, and filters alerts automatically, presenting analysts with prioritized incidents rather than thousands of raw alerts.
- Dwell Time: Faster detection plus faster response equals a smaller window of attacker access.

Q22. How would you design a human-in-the-loop framework for AI security automation?

A human-in-the-loop framework for security automation defines which actions AI can take autonomously and which require human approval before execution. I would design it across three tiers:

- Tier 3: Human-Initiated Only - External breach notification, law enforcement engagement, major system shutdown. AI provides analysis and drafts communications; humans make all decisions.
- Tier 2: AI Recommends, Human Approves - Isolating a user endpoint, suspending a user account, initiating incident escalation. AI prepares the action with full context; a human approves within a defined time window.
- Tier 1: Fully Automated - Blocking a known malicious IP, quarantining a file matching a known malware hash, sending an automated alert notification. High confidence, low blast radius, easily reversed.

The boundary between tiers should be reviewed quarterly as AI confidence levels are validated against real outcomes.

Q23. What is Zero Trust and how does Gen AI enhance its implementation?

Zero Trust is a security model built on the principle of "never trust, always verify" — every access request is authenticated, authorized, and continuously validated regardless of whether it originates inside or outside the network perimeter. Gen AI enhances Zero Trust implementation in three meaningful ways:

- **Continuous Risk Scoring**AI calculates a real-time risk score for every access request based on user identity, device health, location, time, behavior patterns, and resource sensitivity — enabling dynamic access decisions that adapt to current risk rather than static policy rules.
- **Intelligent Step-Up Authentication**AI determines when additional verification is required based on risk context, triggering MFA step-up only when something is genuinely anomalous rather than for every access event, balancing security with user experience.
- **Access Pattern Learning**AI identifies over-permissioned accounts and unused permissions that expand the attack surface unnecessarily, supporting the least-privilege principle at scale across complex environments.

Q24. How has Gen AI changed the phishing threat landscape and what should defenders do differently?

Gen AI has fundamentally broken the traditional heuristics defenders used to identify phishing. Previously, phishing emails could often be identified by poor grammar, generic salutations, and implausible scenarios. Gen AI creates phishing emails that are grammatically perfect, contextually personalized, culturally appropriate, and timely. More importantly, Gen AI enables spear phishing at scale.

Defenders must respond by:

- Moving from content-based phishing detection to behavioral analysis of email metadata and link patterns
- Implementing AI-powered email security tools that can match the sophistication of AI-generated content
- Substantially increasing phishing simulation frequency using AI-generated examples
- Reducing the blast radius of successful phishing through strong MFA and privileged access management

Q25. What are deepfakes and how are they being used in cyberattacks?

Deepfakes are AI-generated synthetic media — video, audio, or images — that realistically impersonate real people by learning and replicating their appearance, voice, and mannerisms from existing content. In cyberattacks, deepfakes are being used in several documented attack scenarios:

- Business Voice Compromise: Attackers call financial staff using a synthetic voice replicating the CEO, authorizing fraudulent wire transfers.
- Video Verification Attacks: Deepfake video is used to pass video-based identity verification for account takeover.
- Social Engineering: Deepfake video messages from "the CISO" instructing staff to take specific actions that benefit attackers.

The defense response requires: implementing verification protocols that don't rely solely on voice or video, deploying deepfake detection tools, and establishing callback procedures for all high-value financial transactions regardless of the channel through which they were requested.

Q26. What is LLM-assisted exploit development and how does it change the threat landscape?

LLM-assisted exploit development refers to attackers using code-capable AI models to understand newly disclosed vulnerabilities and generate working proof-of-concept exploits significantly faster than was previously possible. The traditional exploit development timeline — from CVE publication to weaponized exploit available to attackers — was measured in days to weeks for most vulnerabilities, giving defenders time to patch before widespread exploitation.

AI compresses this timeline toward hours in some cases, because LLMs can read a CVE description, understand the underlying code vulnerability, and generate exploit code without the manual reverse engineering that previously gated this process. This changes the threat landscape by:

- Shrinking the effective patching window
- Increasing the premium on vulnerability prioritization and virtual patching
- Making rapid response to critical CVE disclosures a genuine operational requirement rather than a best practice

Q27. How would you explain the AI threat arms race to a CISO who is skeptical about AI investment in security?

The decision is not whether to invest in AI security — it is whether to invest before or after attackers have used AI to outpace your current defenses.

Attackers are already using Gen AI to write better malware faster, generate more convincing phishing at higher volume, and discover vulnerabilities in your environment more efficiently. Your existing security tools were designed to detect threats generated at human speed and complexity. AI-generated threats operate at machine speed and scale.

The question is not "do we need AI security?" — you need it to maintain the detection capability you already think you have. The question is "which AI security investments are highest priority given our specific threat profile?"

I would then propose starting with the use case where the ROI is clearest: AI-assisted alert triage, which reduces analyst workload while improving detection quality — a measurable, defensible starting point that builds internal confidence for broader AI security investment.

Q28. What ethical challenges arise specifically from AI-powered security monitoring?

AI-powered security monitoring creates genuine ethical tensions that security leaders must navigate deliberately. The core tension is between security effectiveness and employee privacy — comprehensive behavioral monitoring that makes AI anomaly detection effective requires analyzing employee activity at a level of granularity that raises legitimate privacy concerns, particularly in jurisdictions with strong worker privacy protections. Specific challenges include:

- Behavioral profiling that could reveal personal information unrelated to security (health conditions inferred from browsing patterns, political views inferred from communications)
- AI systems that flag certain demographic groups disproportionately due to bias in training data
- The risk of AI-generated behavioral data being misused for performance management or disciplinary purposes beyond its security intent
- Employee consent challenges when monitoring occurs at the infrastructure level without transparent disclosure

Q29. What is the NIST AI Risk Management Framework and how does it apply to cybersecurity AI deployments?

The NIST AI RMF is a voluntary framework providing guidance for managing risks associated with AI systems across their lifecycle. It is organized around four functions:

- Govern: Establish policies, roles, and accountability structures for AI security tools — who is responsible for AI decisions, how AI use is documented, what escalation processes exist when AI underperforms.
- Map: Identify and categorize the AI risks specific to each security tool — for a UEBA system: bias risk, false positive risk, and privacy risk; for an automated response system: the risk of inappropriate automated actions.
- Measure: Define and track metrics for AI performance, bias, and reliability — detection rates, false positive rates, demographic fairness audits.
- Manage: Implement controls to address identified risks and establish processes for continuous monitoring and improvement.

Q30. How does the EU AI Act affect cybersecurity AI deployments?

The EU AI Act classifies AI systems by risk level, and several cybersecurity AI applications fall into categories with significant compliance implications. AI systems used in critical infrastructure — including security systems protecting critical infrastructure — are classified as high-risk and face stringent requirements including conformity assessment before deployment, comprehensive technical documentation, human oversight mechanisms, and high accuracy and robustness standards.

For cybersecurity professionals in EU-regulated organizations or serving EU customers, practical implications include:

- Documenting all AI security tools and their risk classifications
- Implementing human oversight for all high-risk AI security decisions
- Maintaining audit trails of AI decisions for regulatory review
- Ensuring AI security vendors can provide the conformity documentation required for high-risk systems

Q31. How would you handle a situation where an AI security system makes a discriminatory decision — flagging a specific demographic group disproportionately?

This is a multi-step response requiring both immediate action and systematic remediation.

- Immediate: Suspend Autonomous Decision-Making: Suspend the AI system's autonomous decision-making capability for the affected decision type — whether account flagging, access restriction, or incident prioritization — and revert to human-reviewed decisions while the issue is investigated.
- Impact Assessment: Conduct an impact assessment: how many individuals were affected, what decisions were made based on the biased AI output, and what harms resulted.
- Root Cause Investigation: Investigate the root cause: is the bias in the training data, the feature selection, the model architecture, or the threshold settings? Engage technical teams, legal counsel, HR, and potentially external AI ethics expertise.
- Remediate: Retrain the model with debiased data, adjust features that serve as proxies for protected characteristics, and recalibrate decision thresholds.
- Pre-Redeployment Bias Audit: Before redeployment, conduct a comprehensive bias audit across all protected characteristics.
- Document, Communicate & Monitor: Document everything — the discovery, investigation, remediation, and reaudit results. Communicate appropriately to affected individuals and regulators as required. Establish ongoing bias monitoring so this is caught earlier in future deployments.

30-Second Answers

Name three Gen AI-powered security tools.

Microsoft Security Copilot, CrowdStrike Charlotte AI, SentinelOne Purple AI.

What does SOAR stand for?

Security Orchestration, Automation, and Response.

What is dwell time?

The time between initial compromise and detection — AI shortens it significantly.

How does AI help with YARA rules?

It generates YARA detection rules automatically from malware samples.

What is the difference between red team and blue team?

Red team simulates attackers; blue team defends and responds.

What is a GAN and why does it matter for security?

Generative Adversarial Network — attackers use it to generate polymorphic malware and synthetic phishing content.

What does XAI stand for and why does it matter in security?

Explainable AI — security AI must explain its decisions for accountability, compliance, and analyst trust.

What is shadow IT's AI equivalent?

Shadow AI — employees using unapproved AI tools with organizational data, creating security and data leakage risks.

Name two regulatory frameworks relevant to AI in cybersecurity.

NIST AI Risk Management Framework and the EU AI Act.

GSDC
Global Skill Development Council

CERTIFICATION IN GENERATIVE AI IN CYBERSECURITY

ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING OBJECTIVE

- Gain insights into autonomous decision-making processes
- Apply knowledge using ready-to-implement templates
- Demonstrate ability to work with Agentic AI models
- Validate your skills wit

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org