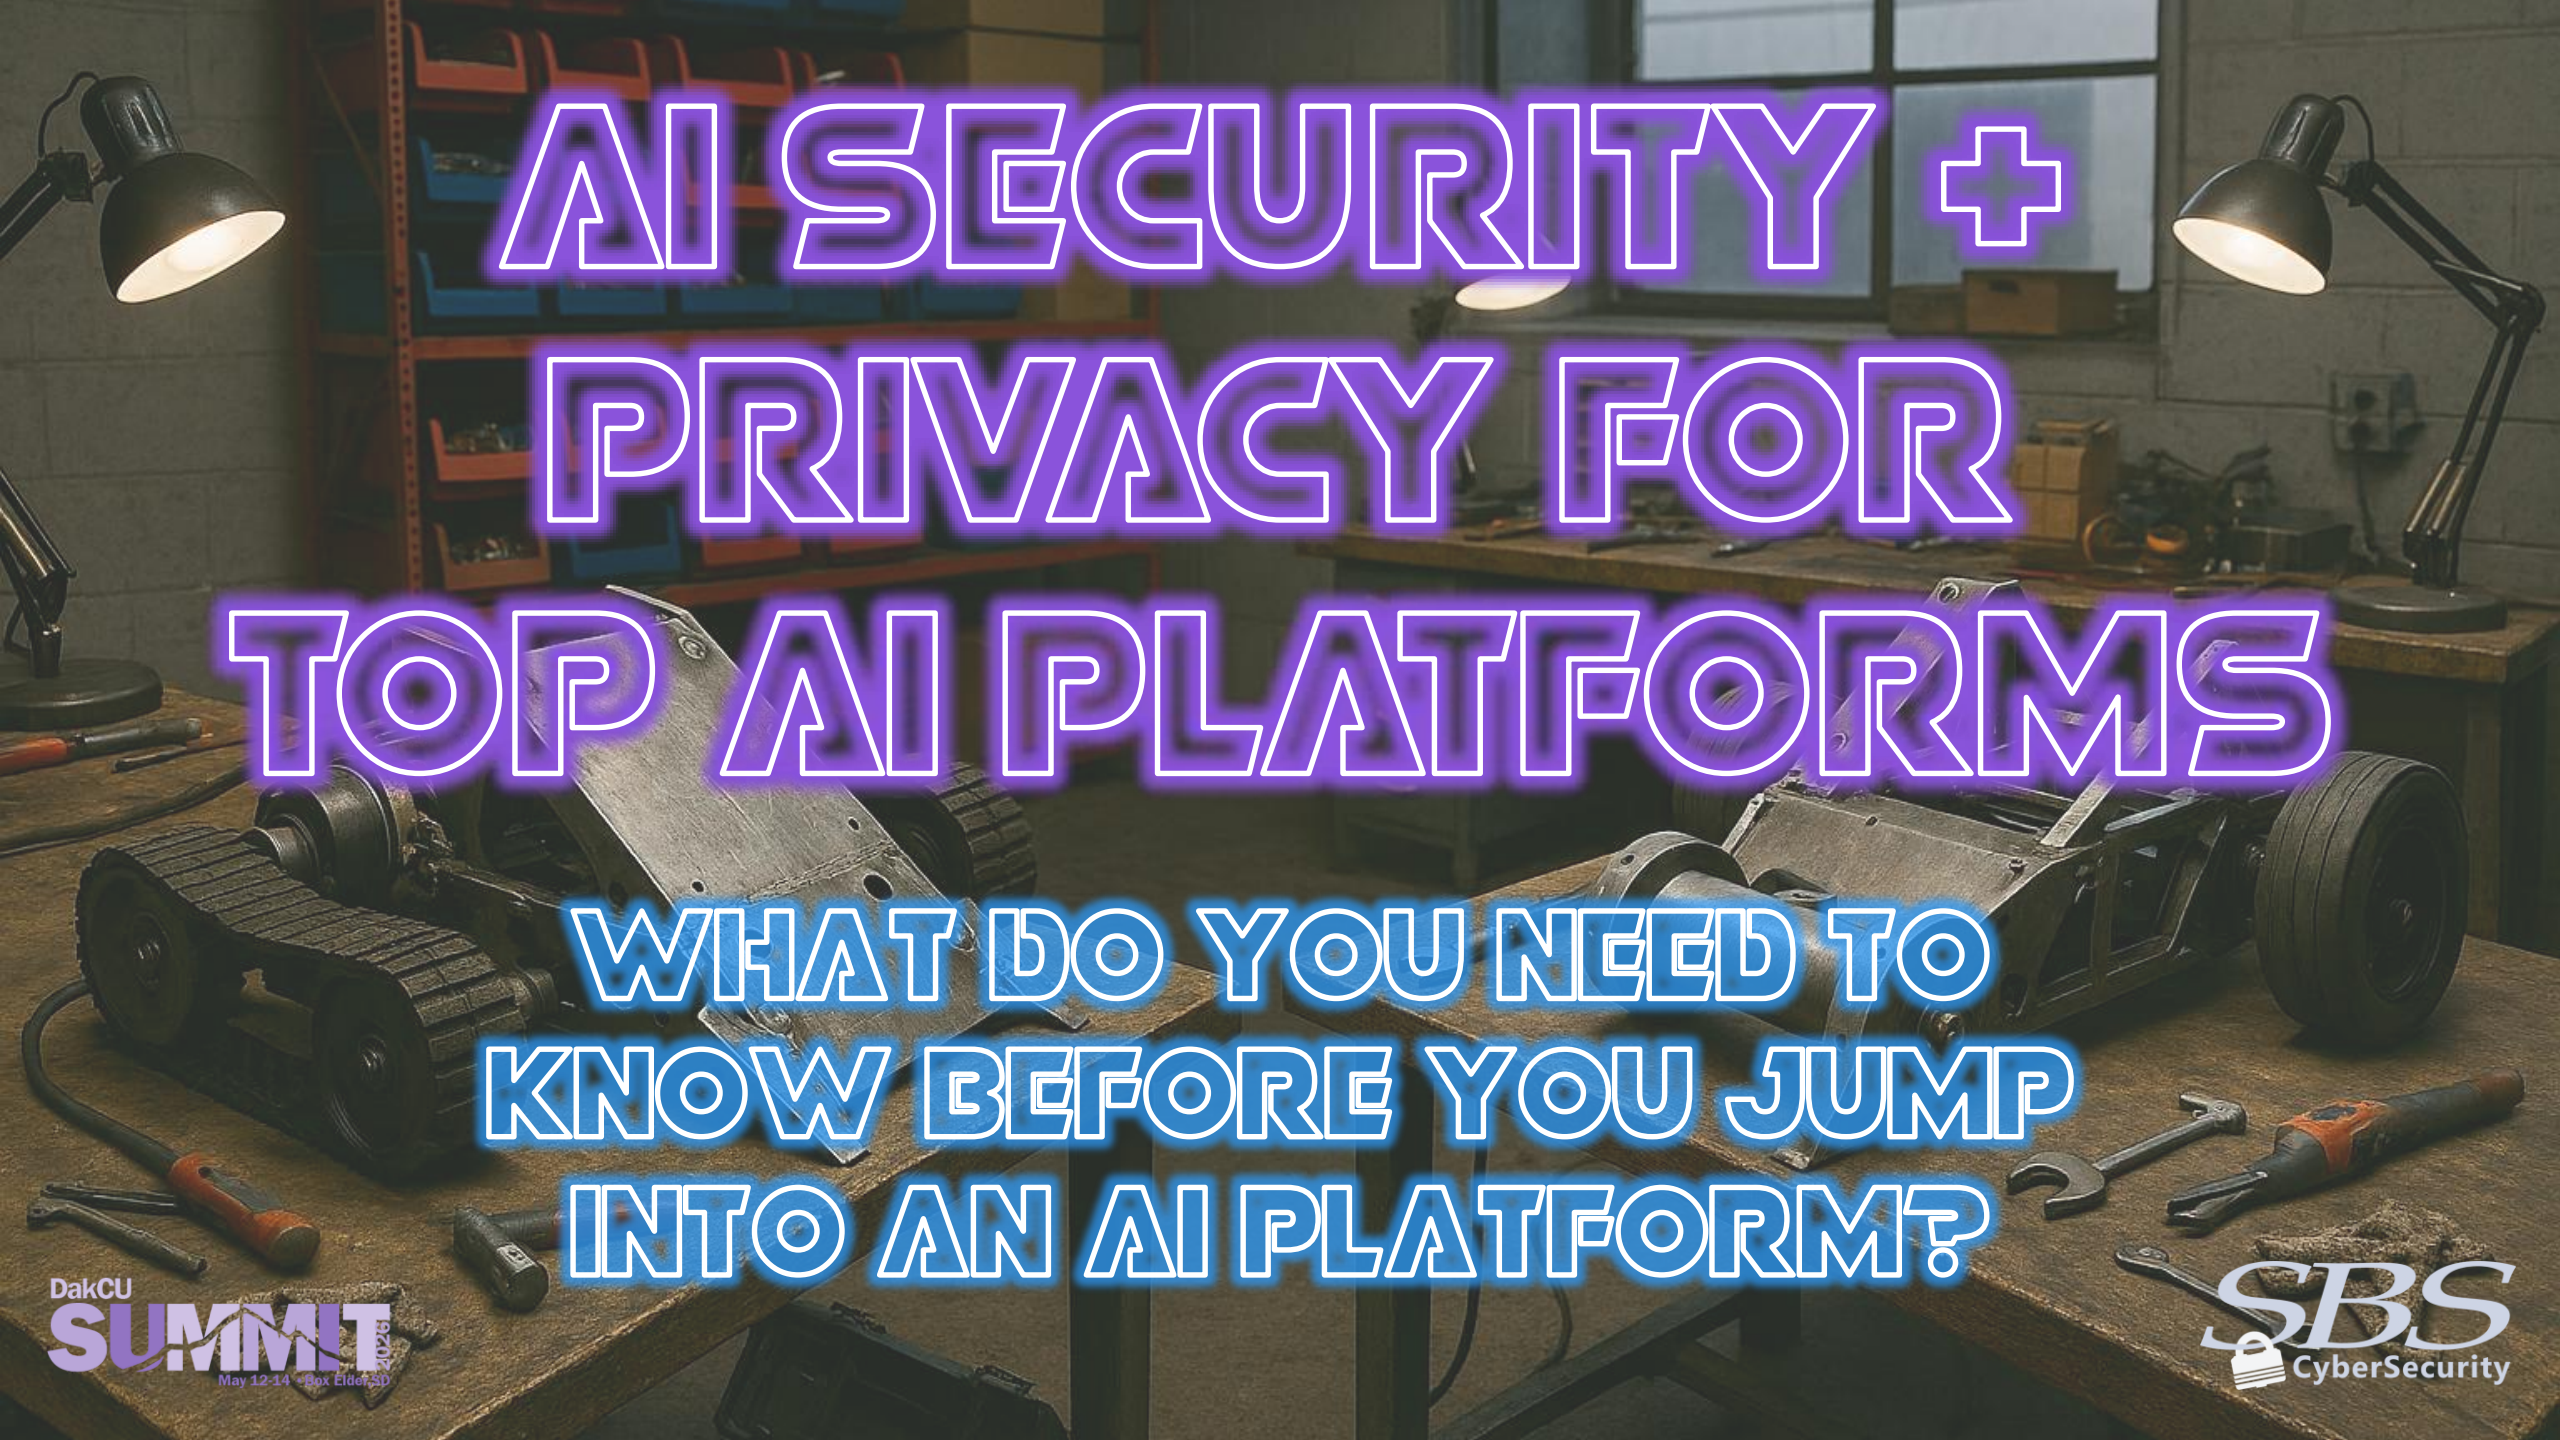
AI SECURITY + PRIVACY FOR TOP AI PLATFORMS
WHAT DO YOU NEED TO KNOW BEFORE YOU JUMP INTO AN AI PLATFORM?
DakCU SUMMIT 2026
May 12-14 • Box Elder, SD
SBS CyberSecurity

JON WALDMAN

- President, Partner, Co-Founder – SBS CyberSecurity
- CISA, CRISC, CPDSE
- Masters of Information Assurance, Dakota State University
- Mission: help you make **empowered** cybersecurity decisions
- Phone: 605-380-8897
- jon@sbscyber.com
- www.sbscyber.com

SBS Institute

- sbsinstitute@sbscyber.com
- 605-269-0909

Follow us on Social:

DOWNLOADS & MORE!

- HEAD TO OUR LANDING PAGE AND DOWNLOAD SOME GOODIES!
- INCLUDING:
- CHANCE TO WIN A FREE SBS INSTITUTE WEBINAR OR MEMBERSHIP!
- TODAY'S SLIDE DECK(S)!
- PRESENTATION SURVEY - WE LOVE FEEDBACK! TELL US HOW WE DID!
- FREE DOWNLOADS - AI POLICY, BLOG POST, INFOGRAPHIC, MORE!
- SIGN UP FOR IN THE WILD!

HTTPS://SBSCYBER.COM/DAKCU

POLLS AND PARTICIPATION

We want to hear from you!

Not required, but encouraged!

Join at

slido.com

Use the code:

#DAKCU2026

What's your role at your institution?

WHAT WE'LL EXPLORE TODAY

- BUILD VS. BUY...?
- AI BOT BATTLE!
  - MICROSOFT COPILOT
  - CHATGPT
  - GOOGLE GEMINI
  - ANTHROPIC CLAUDE
  - PERPLEXITY
  - DEEPSEEK

2 WAYS TO LEVERAGE AI TODAY...

BUILD

VS

BUY

Build Your Own AI - Considerations

BUILD

- Determine Use-Case(s)
- Prepare Infrastructure
- Collect/Curate Data Sources
- Train from Scratch vs. Pre-trained Model
- Train and Validate the Model
- Secure the Environment
- Maintain and Approve

BUY

- Determine Use-Case(s)
- Purchase Licenses
- Secure the Environment
- Validate the Model

If you're leveraging AI today, are you building or buying?

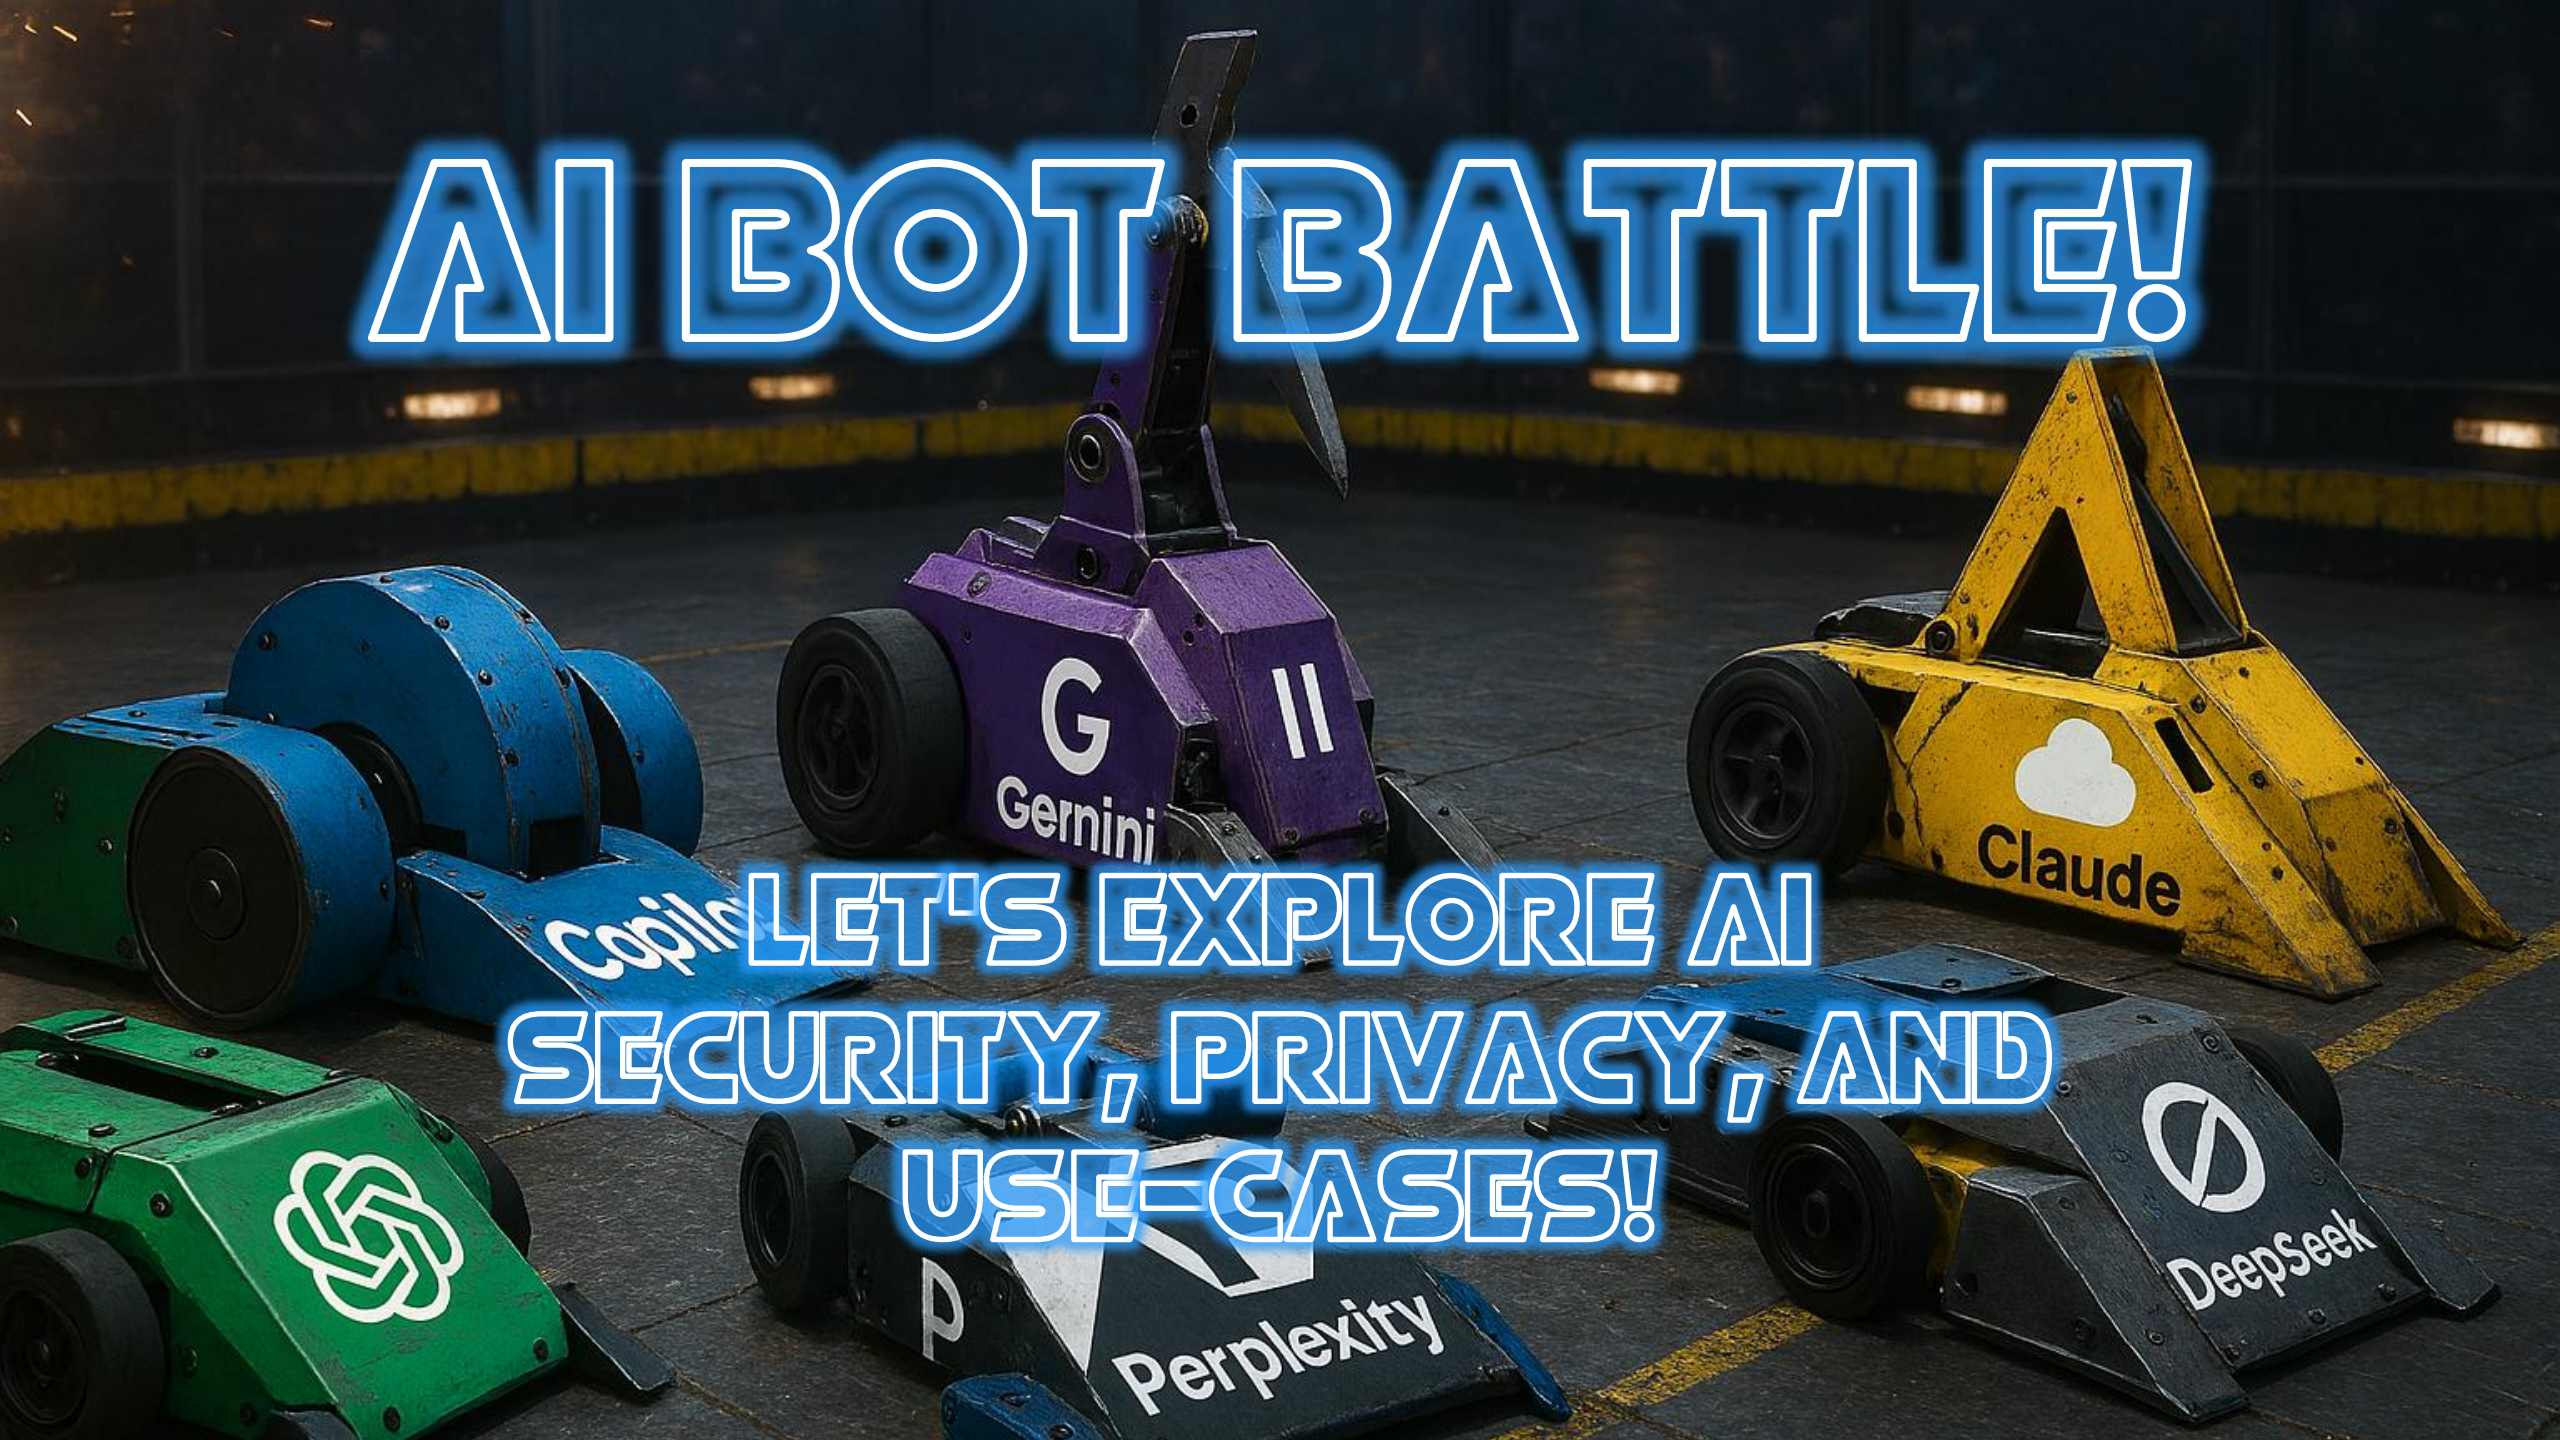
AI BOT BATTLE!
Gernini
Claude
Perplexity
DeepSeek
LET'S EXPLORE AI SECURITY, PRIVACY, AND USE-CASES!

Security and Privacy in AI Models

SBS INSTITUTE

Here's the model for this webinar:

- **Objective:** determine if using <AI Platform> is right for your regulated institution
  - **Question 1**: Is it safe to use this Vendor/Product if I'm a highly regulated financial institution?
  - **Question 2**: Can I trust Vendor/Product to keep my data safe and private?
  - **Question 3**: Why would I choose to use this Vendor/Product?
- **NOTE:** this documentation changes regularly – so this is a point-in-time exercise from 2026 – if you are viewing this later, please use this section as a framework to analyze these statements, as they may have changed

THE BOT BATTLE LINEUP

1. Microsoft Copilot
2. OpenAI ChatGPT
3. Google Gemini
4. Anthropic Claude
5. Perplexity
6. DeepSeek

Which of the following GenAI models have you leveraged at your institution?

Prompt Suggestion

- Here’s the prompt we’re using for this section:
- Please review all available privacy and security documentation from <VENDOR> regarding <AI MODEL> and provide a high-level summary. I'm trying to answer two questions: 1) "Is it safe to use <AI MODEL> if I’m working at a highly regulated financial institution?" and 2) “Can I trust <AI MODEL> and <VENDOR> to keep my data safe and private?

- Mileage may vary depending on which GenAI tool you use.
- Tweak as you see fit.

KEY MESSAGE!

THE BEST SECURITY FEATURES OF MOST SAAS PRODUCTS TODAY ARE PAY TO PLAY

MICROSOFT COPILOT

- **Copilot is the natural first-step for most financial institutions**
- **Question**: Is it safe to use Microsoft Copilot if I'm a highly regulated financial institution?
- **Question**: Can I trust Microsoft Copilot to keep my data safe and private?
- **Question**: Why would I choose to use this Copilot?

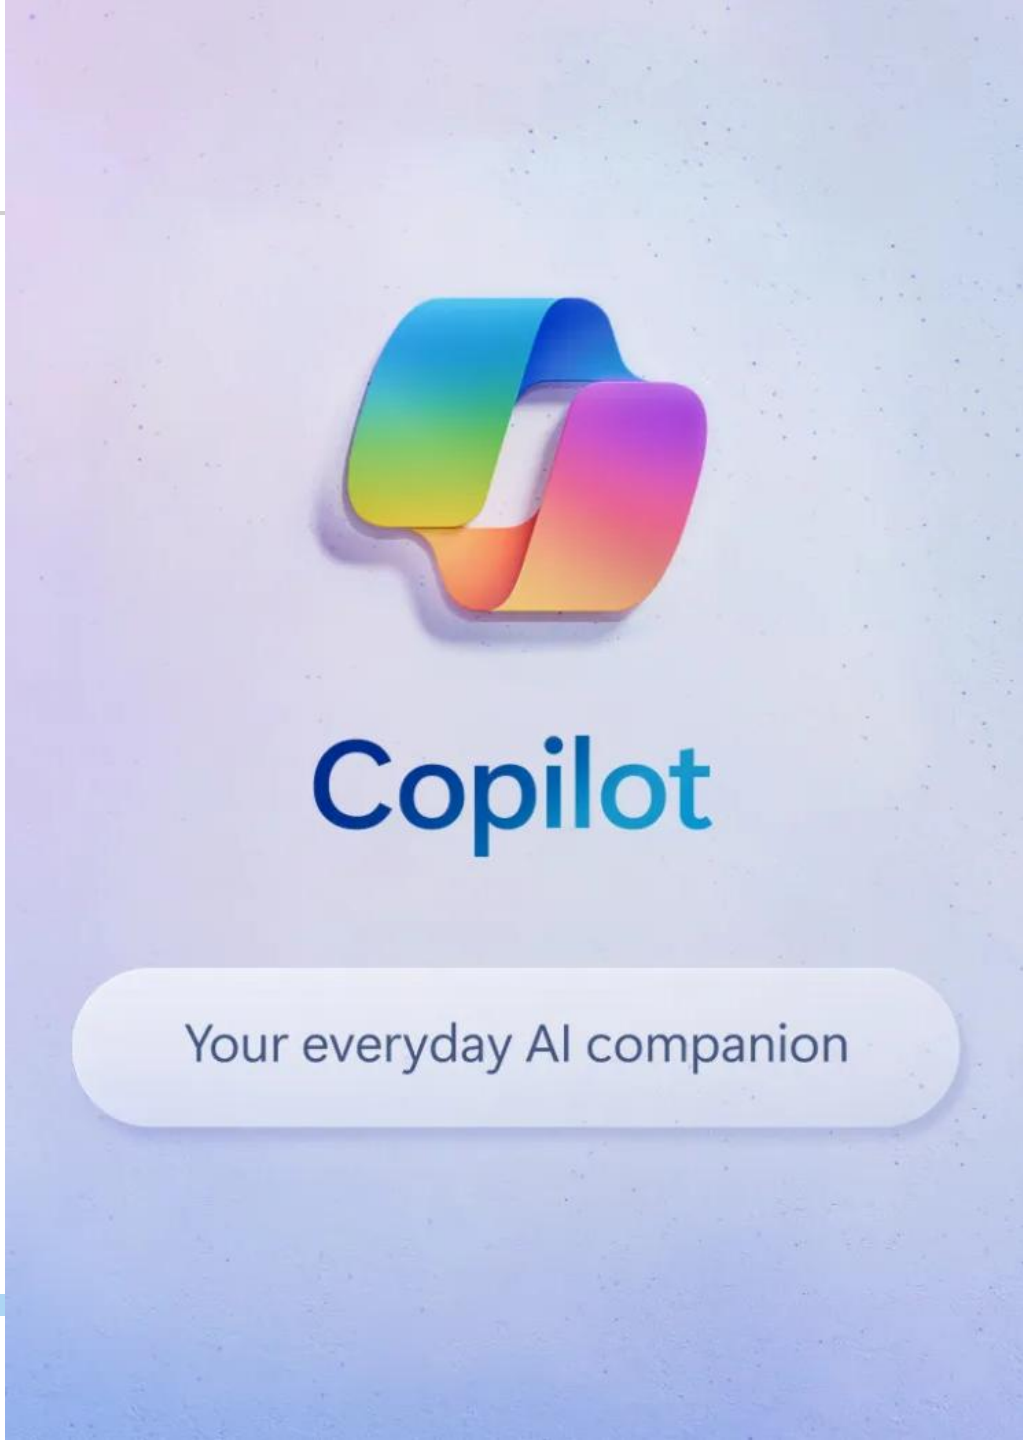

Category	Copilot Chat	Copilot Business (SMB)	Microsoft 365 Copilot (Enterprise add-on)	Copilot Studio
What it is	Free, secure **AI chat** for work. Great for web-grounded answers and quick tasks.	Full-featured **AI assistant for small/midsize businesses** (≤300 users) embedded in Word, Excel, Outlook, Teams, PowerPoint.	Premium **AI assistant for enterprise**—adds Copilot across Microsoft 365 apps with work-data grounding.	Low-code platform to **build custom AI agents** and automate processes; billed with **Copilot Credits** (PAYG or prepaid).
Best for	Everyday help: writing, summarising, web research, light analysis on uploaded files.	Teams that want powerful AI **inside** their apps at SMB pricing.	Organisations needing deep integration with **work data** (emails, files, meetings) and advanced agents.	Creating **custom agents** (e.g., department assistants) that connect to tools, data, and workflows.
Where you use it	Web (m365copilot.com), Teams, Outlook, Edge, Windows app.	Directly in Microsoft 365 apps: Word, Excel, PowerPoint, Outlook, Teams, plus the Copilot app (Chat, Pages, Search, Agents).	Same in-app experiences across Microsoft 365, plus Copilot Chat, Pages, Search, Notebooks, Agents.	In Copilot Studio (web), publishing to channels (e.g., SharePoint, Teams, web) and integrating via connectors.
Data it uses ("grounding")	**Web data**. Can use **uploaded files** you provide in chat; doesn't browse your tenant by default.	Your **work data** (Microsoft Graph: emails, files, chats, calendars) with your permissions.	Your **work data** across Microsoft 365 + partner data via Graph connectors; personalisation with Work IQ.	Whatever you connect: internal/external sources, plugins, tools; consumption metered as **Copilot Credits**.
Security & privacy	**Enterprise Data Protection**; admin controls for safe business use.	Same enterprise protection; respects tenant permissions and policies.	Same, plus admin governance (Copilot Control System, Purview, agent management).	Operates within Power Platform/M365 governance; licensing & usage tracked via **Copilot Credits**.
Agents	Basic web-grounded agents; **work-grounded agents** available **metered** via PAYG.	**Included** agents (Researcher, Analyst, etc.) for SMB scenarios inside apps and chat.	**Full** agent set (create, discover, deep reasoning like Researcher/Analyst) grounded in work data.	**Build custom agents**; design conversations, connect tools, publish and operate; pay by **Copilot Credits**.
Cost	**Included** with eligible Microsoft 365 subscriptions.	**USD $21/user/month** (SMB pricing).	**USD $30/user/month** add-on for enterprise (requires qualifying base plan).	**PAYG** via Azure or **prepaid packs**; price based on **Copilot Credits** consumption.
License prerequisites	Available to commercial customers with eligible Microsoft 365 subscriptions (no extra license).	Requires **Microsoft 365 Business** plan; limited to ≤300 seats.	Requires a **qualifying Microsoft 365 or Office 365 plan** (E3/E5/F-series, Business plans, Apps).	Needs a Copilot Studio **user license** and a tenant **Copilot Credits** subscription (or PAYG).
Typical limitations	Web-only grounding unless you upload files; lighter capabilities vs paid Copilot in apps.	Seat cap (≤300); feature breadth targeted for SMB simplicity and price.	Additional cost per user; depends on data readiness and governance.	Variable spend (credits); requires design/ops skills to build and manage agents.

https://tracyvanderschyff.com/2026/01/07/copilot-chat-vs-business-vs-enterprise-vs-studio-clear-simple-decoded/

Is Copilot safe to use?

- **Yes—when deployed via Microsoft 365 with enterprise configurations.**
- Microsoft Copilot is designed to meet the stringent requirements of regulated industries, including financial services.
- Copilot can be considered safe, but with important caveats and responsibilities on your part.
- NOTE: Authentication and Authorization to Copilot is handled through EntraID – meaning that permissions are inherited by the user via policy or Conditional Access

Is Copilot safe to use?

✅ Regulatory Compliance

- Microsoft 365 Copilot is designed to comply with a wide range of global privacy laws and standards, including GDPR, the EU Data Boundary, ISO/IEC 27001, and ISO/IEC 27018. Microsoft actively works to ensure its services meet regulatory requirements, including those relevant to financial services (e.g., SEC, FINRA, CFTC rules).
- **Immutable storage and retention policies**: Available via Azure Blob Storage and Microsoft Purview.

✅ Enterprise Data Protection

- **Encryption at rest and in transit** using FIPS 140-2–compliant technologies.
- **Data isolation between tenants** and support for Double Key Encryption (DKE).
- **Copilot respects existing Microsoft 365 permissions, sensitivity labels, and retention policies**.

✅ Deployment Best Practices

- **Zero Trust architecture**: Microsoft recommends applying least privilege, MFA, device compliance, and threat protection before assigning Copilot roles.
- **Phased rollout**: Start with pilot groups, monitor usage, and refine policies before full deployment.

Can I trust Copilot and Microsoft?

- **Yes—with enterprise-grade controls and governance.** Microsoft Copilot is built on a foundation of responsible AI, privacy-by-design, and layered security.

🔐 Privacy Commitments

- **Microsoft acts as a data processor** under the Data Protection Addendum (DPA) and Product Terms.
- **Prompts, responses, and Microsoft Graph data are not used to train foundation models**.
- **Web queries via Bing are anonymized and not used for training**.

🛡️ Security Architecture

- **Defense-in-depth strategy**: Includes secure engineering, threat intelligence, red teaming, and containment-by-design.
- **Prompt injection defenses**: Microsoft uses classifiers, markdown sanitization, and session hardening to block jailbreaks and malicious inputs.
- **Audit logs and eDiscovery**: Available via Microsoft Purview for monitoring and compliance.

⚠️ Risk Mitigation

- **Oversharing prevention**: Copilot only surfaces data users already have access to; admins should audit permissions regularly.

CoPilot Common Use Cases

- Automated Financial Reporting
- Budgeting & Forecasting
- Compliance & Risk Monitoring
- Member Support
- Document & Contract Drafting/Review
- Audit and Risk Operations Support
- Meeting and Email Summarization
- AI Agent Creation for Repeatable, Internal Workflows

USE COPILOT IF YOU…

- Are a Microsoft 365 shop and want to harness the power of GenAI
- Copilot is the natural and “safest” first step for most financial institutions…
- And you’ve already done your homework on Microsoft as a vendor

Is your Institution using Copilot?

OPENAI CHATGPT

- **ChatGPT is the most popular AI Platform in the world today**
- **Question**: Is it safe to use ChatGPT if I'm a highly regulated financial institution?
- **Question**: Can I trust ChatGPT to keep my data safe and private?
- **Question**: Why would I choose to use this ChatGPT?

Free

Intelligence f

$0 / mon

Limited m

Limited ar

nited de

Limited

Have an existing

Business

A secure, collaborative workspace for startups and growing businesses

Annual billing | Monthly billing

$20 / user / month

Try for free >

Everything in Plus and:

- Assign standard or usage-ba… eam's needs. Learn more
- Unlimited GPT-5.5 me… king, and access to GPT-5.5 Pr…
- 60+ apps that bring … gle Drive, SharePoint, GitHub, A…
- A secure, dedicated wo… SSO, and MFA
- Support for compliance with … aws. Aligned with CSA STAR and SOC 2 Type 2
- Business features like apps, data analysis, record mode, canvas, shared projects, and custom workspace GPTs
- Encryption at rest and in transit, and no training on your business data by default. Learn more
- Includes access to Codex for reasoning and taking action across your documents, tools, and codebases

Unlimited subject to abuse guardrails. Learn more

Enterprise

Enterprise-grade AI, security, and support at scale

Contact sales >

Everything in Business and:

- Expanded context window that s… files
- Enterprise-level security and… ytics, domain verification, and ro…
- Advanced data privacy w… rest and in transit, and no tra… e
- Support for data residen… ors
- 24/7 priority support, SLAs, (eligible customers)
- Invoicing and billing, volume discou…

ty

e uploads

ge creation

and agent

ontext

, and custom

features

rdrails. Learn

Is it safe to use ChatGPT?

- **Yes—with the right deployment model.** GPT-5, especially when used via **ChatGPT Enterprise**, **ChatGPT Business**, or the **OpenAI API**, is designed to meet the needs of regulated industries like finance. Controls include:
  - ✅ **Enterprise-grade security**: AES-256 encryption at rest, TLS 1.2+ in transit.
  - ✅ **SOC 2 Type 2**, **ISO/IEC 27001**, **27017**, **27018**, and **27701** certifications.
  - ✅ **No training on business data by default**: Inputs and outputs from enterprise products are not used to train models unless explicitly opted in.
  - ✅ **Data residency options**: Available to support local data sovereignty requirements.
  - ✅ **Support for HIPAA compliance**: Business Associate Agreements (BAAs) are available for eligible customers.
- However, **default consumer versions (e.g., ChatGPT Free or Plus)** are **not recommended** for regulated environments due to data retention and training policies.

Can I trust ChatGPT?

Yes—with enterprise-grade deployments and proper configuration.
OpenAI has implemented a **multi-layered security and privacy framework**:

🔐 Privacy Controls

- **You own your data**: Inputs and outputs are yours, and OpenAI only retains rights necessary to provide services.
- **Custom retention policies**: Admins can control how long data is stored; deleted data is purged within 30 days unless legally required.
- **Zero Data Retention (ZDR)**: Available for eligible API endpoints.

🛡️ Security Measures

- **External penetration testing** and a **Bug Bounty Program** to identify vulnerabilities.
- **Strict access controls** and 24/7 incident response coverage.
- **Trust Portal**: Offers transparency into certifications, audits, and system architecture.

🤖 AI Safety Innovations

- **Safe-Completions Training**: GPT-5 uses a new safety paradigm that avoids binary refusal and instead provides helpful, safe responses to ambiguous or dual-use prompts.
- **Red-teaming and adversarial testing**: Conducted internally and with partners like Microsoft and the UK AI Safety Institute.

ChatGPT Common Use Cases

SBS INSTITUTE

- 🚀 Access to (typically) the most advanced models
- 🔍 Risk Modeling & Fraud Detection
- 📄 Regulatory Reporting Automation
- 💼 Personalized Financial Advice
- 🤝 Enhanced Member Support
- 👩‍💻 Developer Enablement or Support Agents via API

USE CHATGPT IF YOU...

- Don’t use Copilot and want the most up-to-date, functional AI models
- Are creating custom GPTs
- Leveraging an AI model via API

GOOGLE GEMINI

- **Google is also going all-in on AI, integrating Gemini into phones, home IoT devices, Workspace, and more…**
- **Question**: Is it safe to use Gemini if I'm a highly regulated financial institution?
- **Question**: Can I trust Gemini to keep my data safe and private?
- **Question**: Why would I choose to use this Gemini?

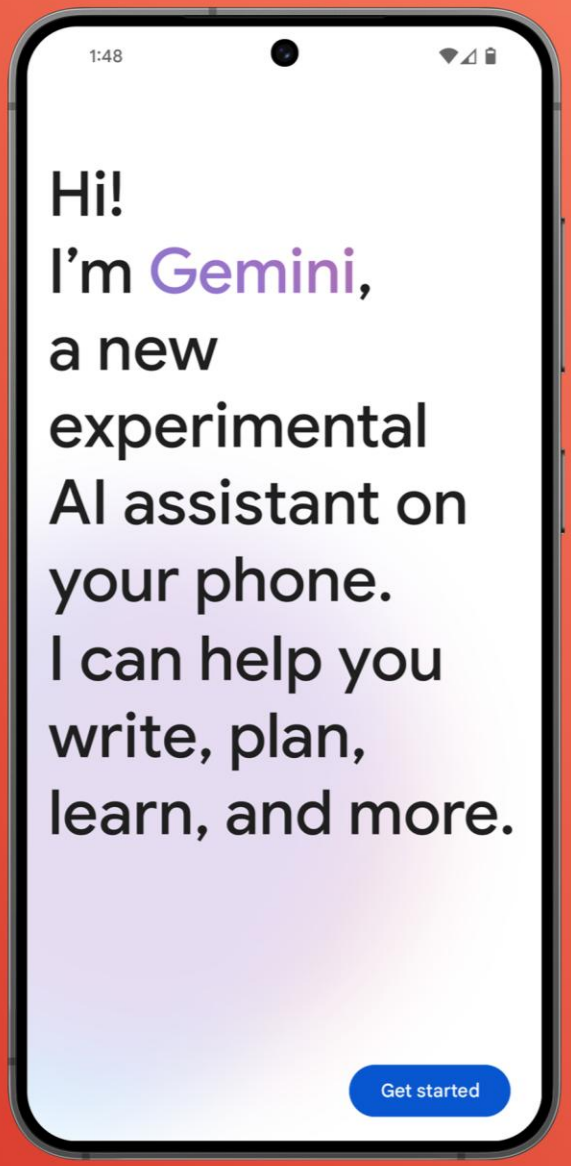

m/lp/business/

Annual (Save 16% with one-year commitment)

Starter

$7

per user / month

Get started

Starter includes:

- 30 GB pooled storage per user*
- Custom business email, you@your-company.com
- Gemini AI assistant in Gmail
- Chat with AI in the Gemini app
- Video meetings, 100 participants
- Security and management controls

Standard

$14

per user / month

Get started

All of Starter, and:

- 2 TB 65x more than
- Custom b you@your- custom la
- Gemini AI Docs, Mee
- AI research a (NotebookLM)
- Chat with AI in the Ge create your team of AI experts
- Video meetings with recording, noise cancellation, 150 participants
- Appointment booking pages
- eSignature with Docs and PDFs
- Google Workspace Migrate tool for data migration

Plus

$22

per user / month

Get started

All of Standard, and:

- han Standard*
- iness email
- with attendance rticipants
- archive, and
- DAP
- Advanced endpoint management
- Enhanced security and management controls

Enterprise

Contact sales for pricing

Contact sales

All features mentioned, and:

- 5 TB or upgrade for more*
- Custom business email + S/MIME encryption
- Video meetings with in-domain live streaming, 1000 participants
- Data Loss Prevention (DLP)
- Context-Aware Access (CAA)
- Enterprise data regions
- Cloud Identity Premium
- Enterprise endpoint management
- AI Classification for Google Drive
- Assured controls available as an add-on
- Enhanced Support with faster response times for critical issues

Is it safe to use Gemini?

Yes—if deployed via Google Cloud or Workspace with enterprise controls.
Gemini is designed to meet the needs of regulated industries, including finance, through robust security, compliance, and governance features:

✅ **Enterprise-Grade Security**

- **Encryption**: TLS for data in transit and AES-256 for data at rest.
- **Identity & Access Management (IAM)**: Role-based access, SSO, and 2-step verification.
- **VPC Service Controls**: Prevents data exfiltration and enforces perimeter security.
- **Client-Side Encryption (CSE)**: Ensures even Google cannot access sensitive data.

✅ **Regulatory Compliance**

- **Supports GDPR, HIPAA, FedRAMP High, and ISO/SOC certifications**.
- **Data residency and governance tools**: Admins can control where data is stored and how it's accessed.

✅ **Enterprise Use Cases in Finance**

- Fraud detection, risk analysis, compliance automation, and secure document summarization.

Important: Consumer-facing Gemini apps (e.g., mobile/web) are **not recommended** for regulated environments due to limited admin controls and broader data usage policies.

Can I trust Gemini and Google?

Yes—with enterprise deployments and proper configuration. Google has implemented a multi-layered privacy and security framework for Gemini:

🔐 Privacy Practices

- **No training on customer data by default**: Prompts and responses are not used to train models unless explicitly opted in.
- **Auto-delete settings**: Default retention is 18 months, configurable by users.
- **Granular admin controls**: Admins can manage access, retention, and export of Gemini data.

🛡️ Security Innovations

- **Model Hardening**: Gemini 3 is trained to resist indirect prompt injection attacks.
- **Automated Red Teaming (ART)**: Google continuously tests Gemini against evolving threats.
- **Audit Logs & Vault Integration**: Full visibility into Gemini usage and access for compliance and eDiscovery.

⚠️ Risks to Consider

- **Consumer Gemini apps** may retain data longer and lack enterprise-grade controls.
- **Adaptive threats** like indirect prompt injection require ongoing vigilance, even with hardened models.

Gemini Common Use Cases

- Workspace Integrations (like M365)
- Financial Document Search & Synthesis
- Enhanced Virtual Assistants (Gems)
- Capital Markets Research
- Regulatory Code Change Consultant
- Personalized Financial Recommendations

USE GEMINI IF YOU...

- Are a Google Workspace institution…

OR

- For personal use…

ANTHROPIC CLAUDE

- **Claude just recently launched it's "Claude for Financial Services" model…**
- **Question**: Is it safe to use Claude if I'm a highly regulated financial institution?
- **Question**: Can I trust Claude to keep my data safe and private?
- **Question**: Why would I choose to use this Claude?

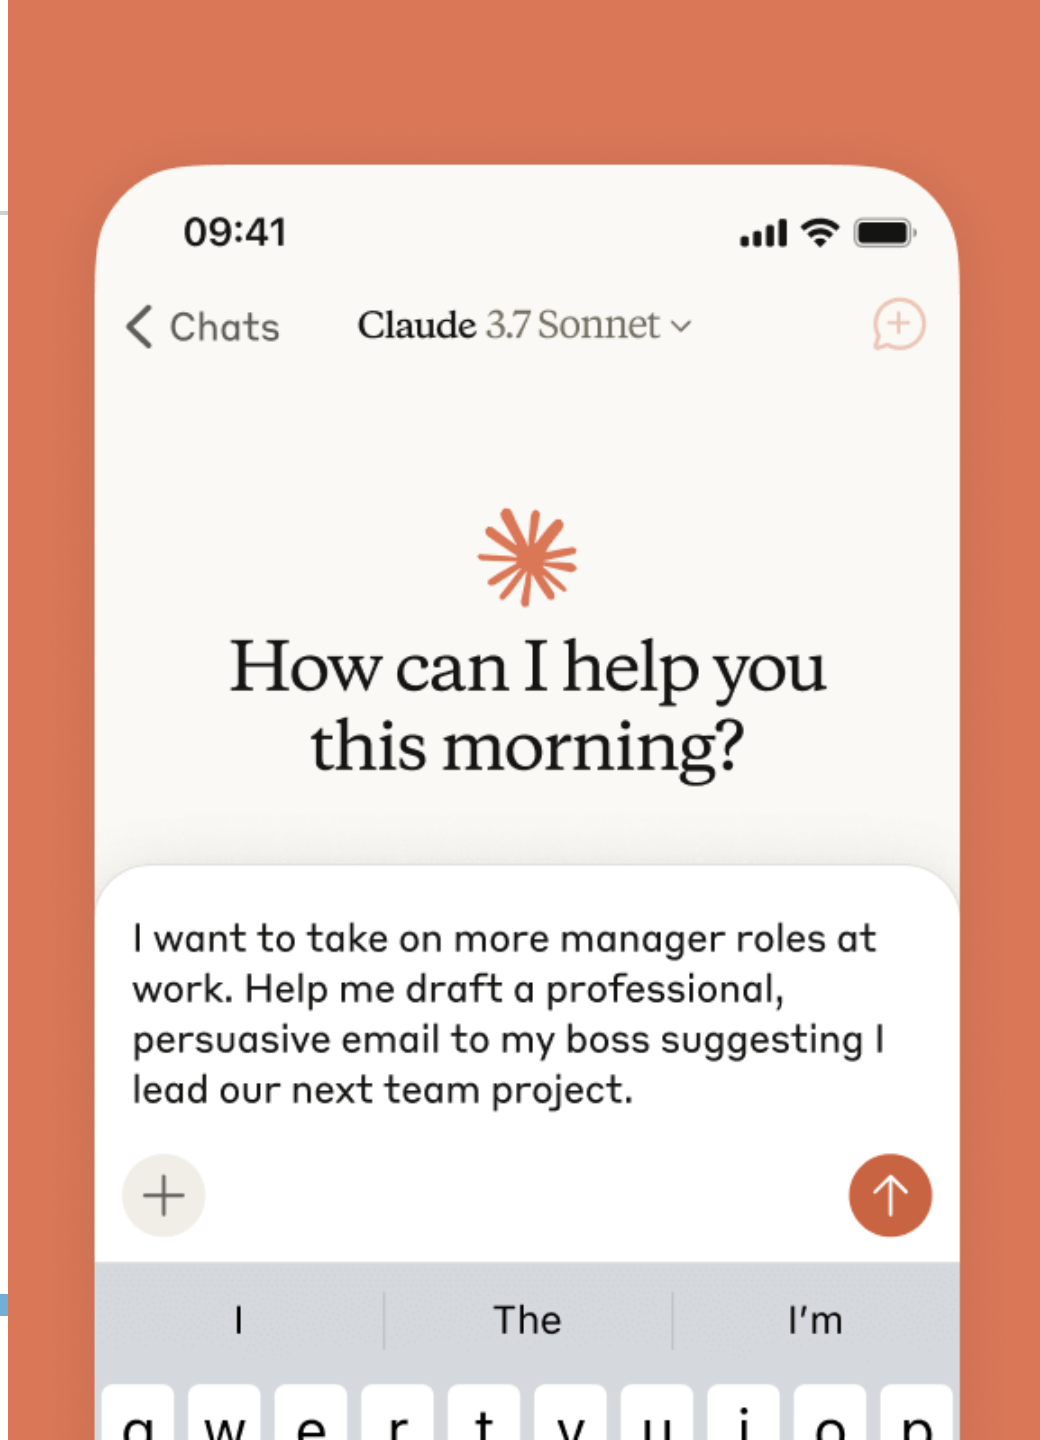

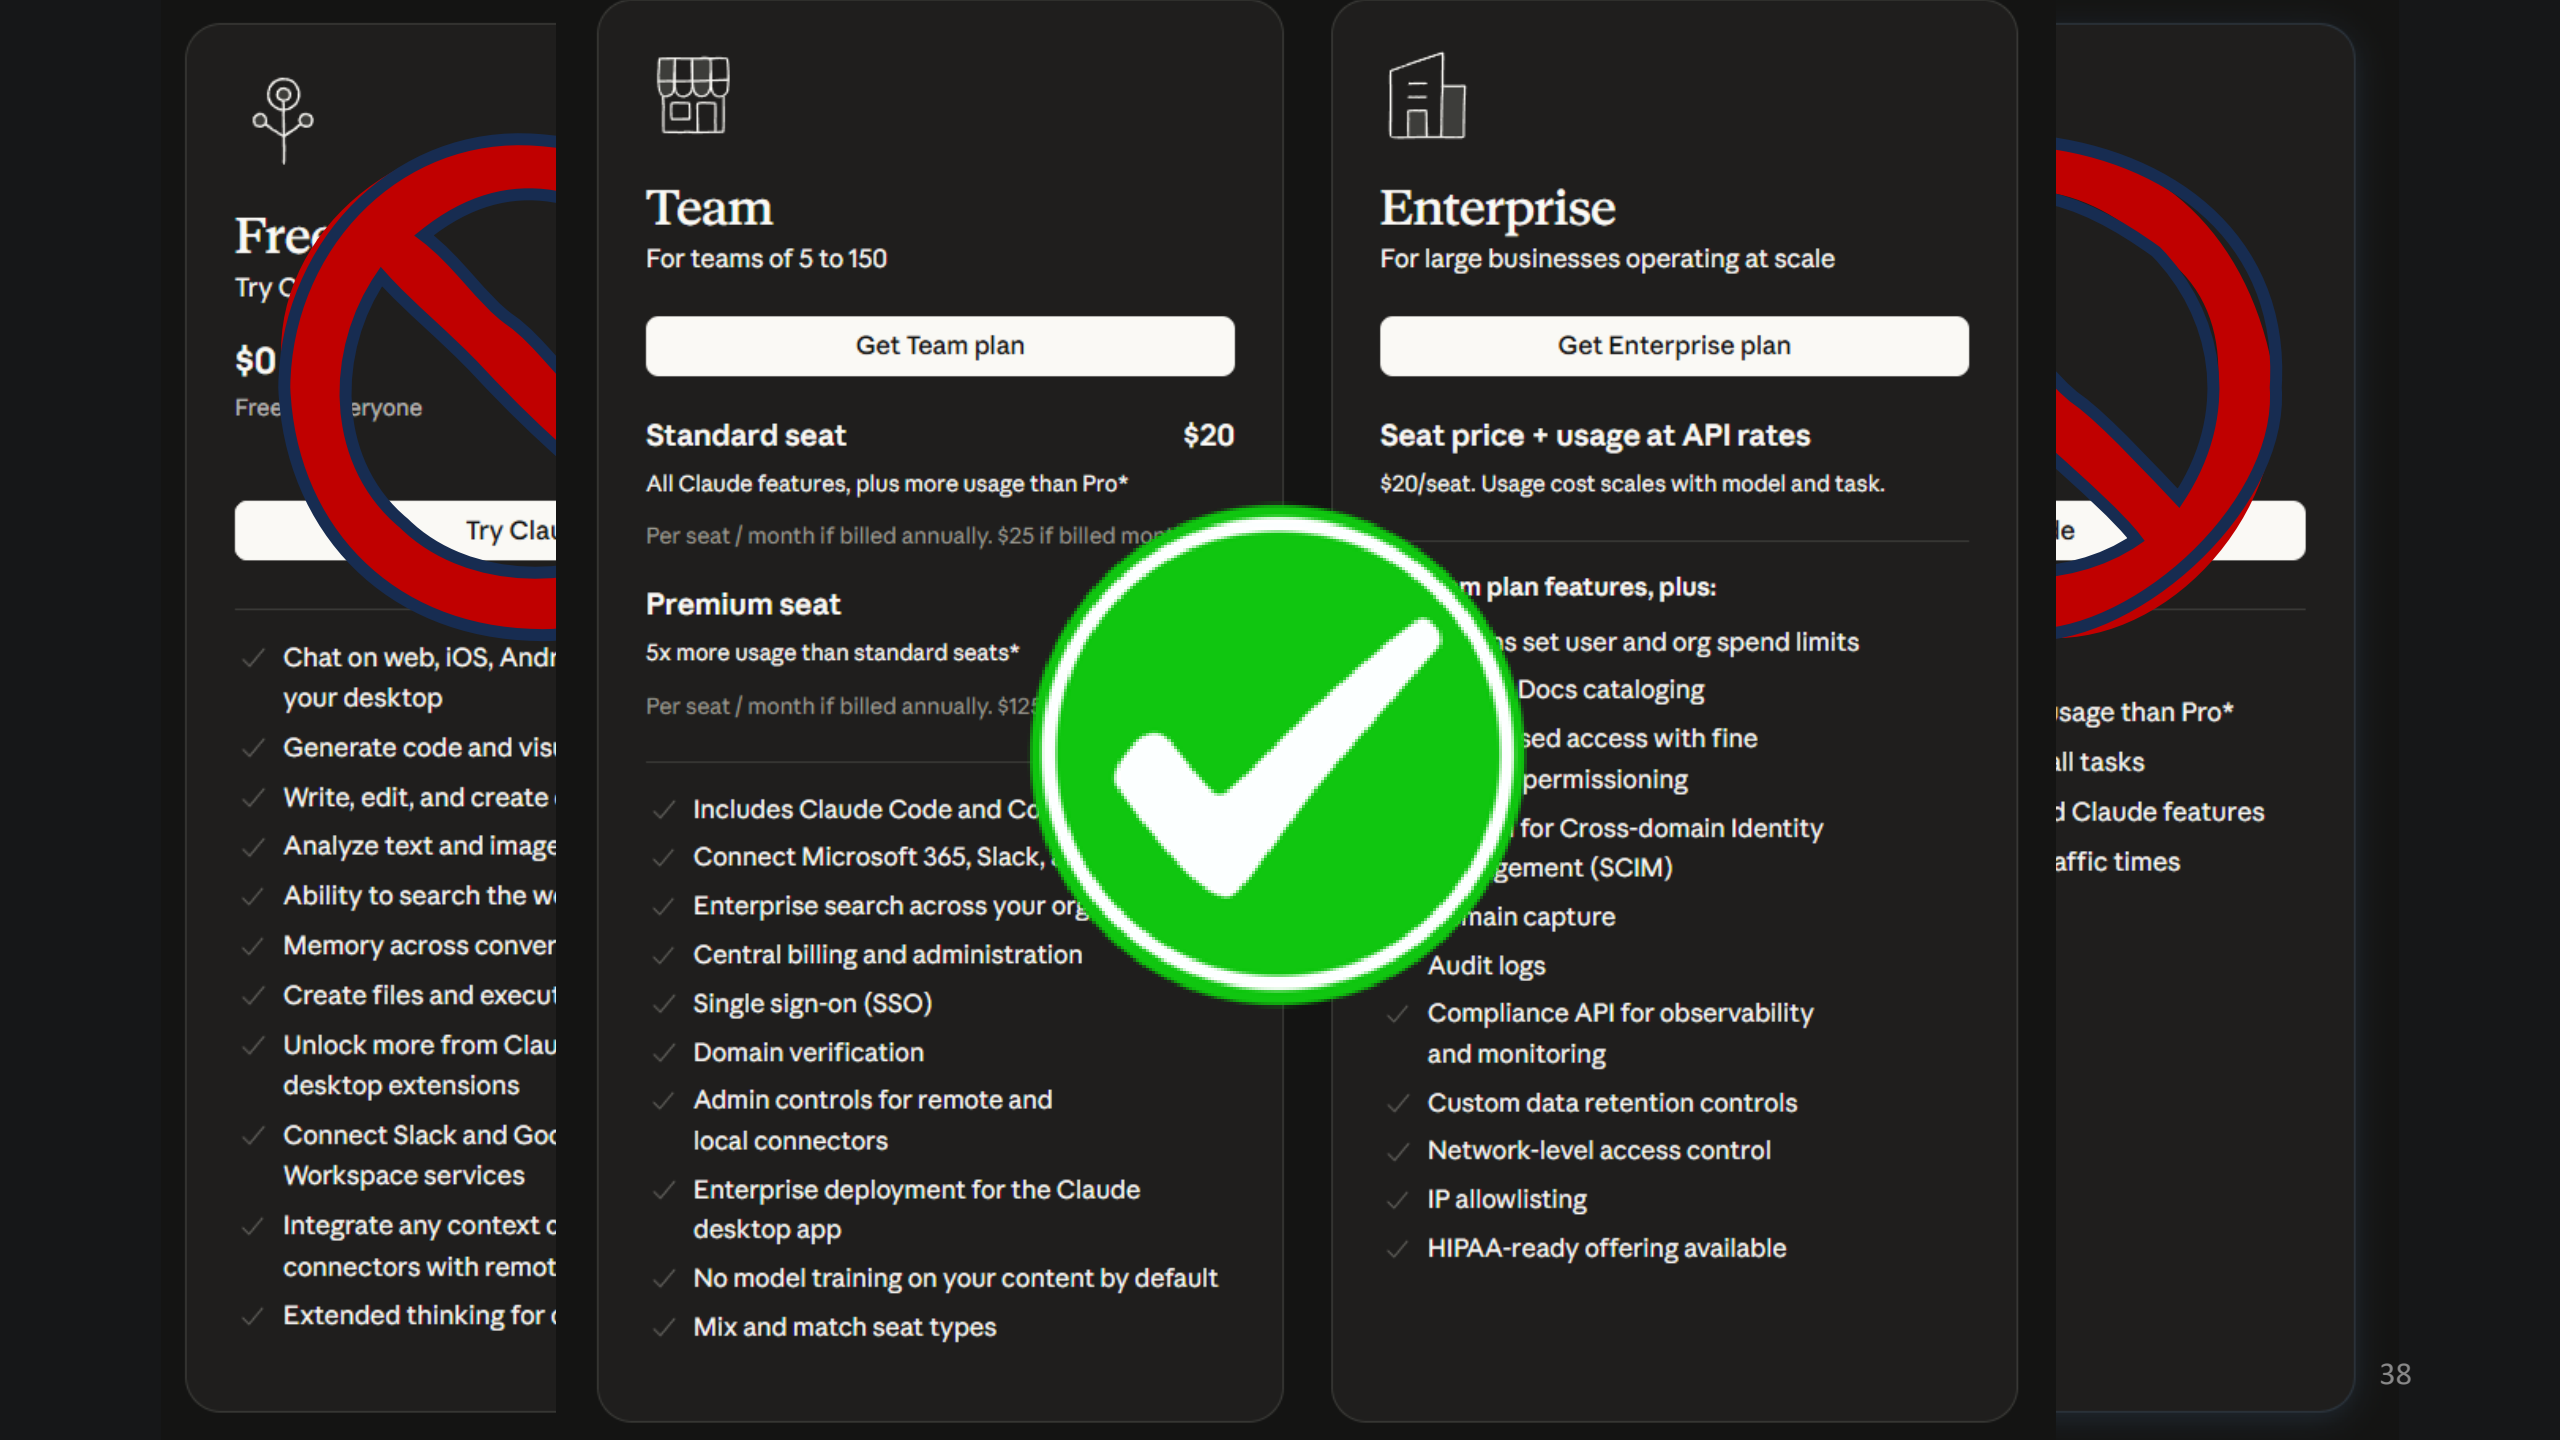
Team
For teams of 5 to 150
Get Team plan
Standard seat
$20
All Claude features, plus more usage than Pro*
Premium seat
5x more usage than standard seats*
Includes Claude Code and Co
Connect Microsoft 365, Slack,
Enterprise search across your org
Central billing and administration
Single sign-on (SSO)
Domain verification
Admin controls for remote and local connectors
Enterprise deployment for the Claude desktop app
No model training on your content by default
Mix and match seat types
Enterprise
For large businesses operating at scale
Get Enterprise plan
Seat price + usage at API rates
$20/seat. Usage cost scales with model and task.
Audit logs
Compliance API for observability and monitoring
Custom data retention controls
Network-level access control
IP allowlisting
HIPAA-ready offering available

Is it safe to use Claude?

Yes—especially with Claude for Financial Services and enterprise-grade deployments. Anthropic has made significant investments in tailoring Claude for high-trust sectors like finance.

✅ Claude for Financial Services (Claude FS – custom pricing)

- **Industry-specific solution** designed for analysts, underwriters, and portfolio managers.
- **Integrations** with Morningstar, S&P Global, Databricks, and Snowflake for unified data access.
- **Expanded context window** for complex modeling and due diligence.
- **Compliance automation tools** like CRaiG (Compliance Requirements Generator).
- **Client data is never used for training**, a critical requirement for financial institutions.

✅ Enterprise Security Features

- **SOC 2 Type 2 and ISO 27001 certifications**.
- **Permission-based architecture**: Claude Code requires explicit approval for sensitive actions.
- **Zero-Data-Retention (ZDR) mode**: Ensures logs are discarded immediately after abuse checks.
- **Custom retention policies**: API logs default to 7 days, with options for longer retention via DPA.

✅ Deployment Options

- Claude is available via **Anthropic's API**, **Claude Enterprise**, and **third-party platforms** like AWS Bedrock and Google Vertex AI, all of which support enterprise-grade controls.

Can I trust Claude?

Yes—with enterprise configurations and proper governance. Anthropic's approach to privacy and safety is comprehensive and transparent.

🔐 Privacy Practices

- **No training on user data by default**: Training requires explicit opt-in, which Anthropic does not currently request.
- **Granular retention controls**: Users can delete chats, and enterprise customers can configure retention windows.
- **Memory features** (for personalization) are opt-in and user-controlled.

🛡️ Security Architecture

- **Input sanitization and command blocklists** to prevent prompt injection.
- **Secure credential storage**, network request approval, and isolated context windows.
- **Real-time enforcement** using classifiers to detect and block harmful content.

🧠 AI Safety & Governance

- **Unified Harm Framework**: Evaluates risks across physical, psychological, economic, societal, and autonomy dimensions.
- **Policy Vulnerability Testing**: Collaborations with experts to stress-test Claude's responses.
- **Usage Policy**: Updated regularly to reflect evolving risks, with specific guidance for high-risk domains like finance.

Claude Common Use Cases

- Large Document Handling
- Developer Workflows (with Claude Code)
- Safer Agent Execution (guardrails)
- Investment Research & Analysis
- Financial Modeling & Forecasting
- Regulatory Compliance Automation
- Underwriting & Risk Assessment
- Document Intelligence & Reporting

USE CLAUDE IF YOU...

- Are processing lots of large (hundreds of pages) documents
- Want the strongest privacy protections and the “strongest” ethical and responsible AI usage
- Want to use Claude Code as your own Dev team

Do you do any software development at your institution?

PERPLEXITY

- **Perplexity is best known for its research and citation capabilities…**
- **Question**: Is it safe to use Perplexity if I'm a highly regulated financial institution?
- **Question**: Can I trust Perplexity to keep my data safe and private?
- **Question**: Why would I choose to use Perplexity?

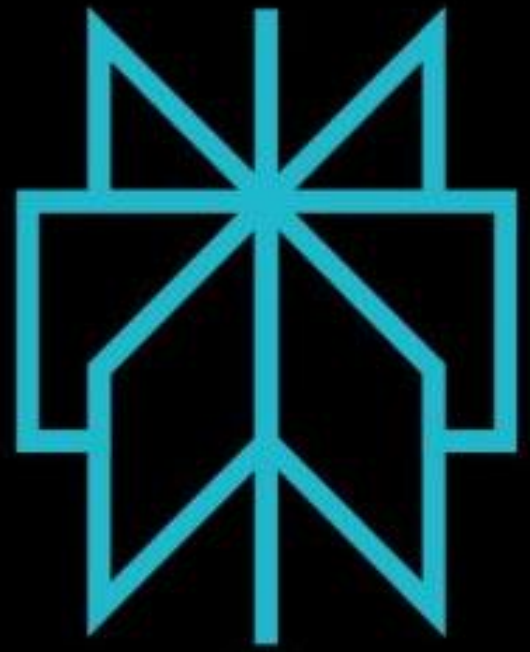

Features	Pro	Enterprise Pro	Enterprise Max
	Learn more	Get started	Get started
Admin Controls + Security			
Single sign on + SCIM*		Integrate with your identity provider for SSO + SCIM	Integrate with your identity provider for SSO + SCIM
User management		Admins can add, remove, and set permissions across team members	Admins can add, remove, and set permissions across team members
Insights dashboard*		Analyze your usage trends and track adoption	Analyze your usage trends and track adoption
Data retention options*		Ability to auto-delete files within 1 day	Ability to auto-delete files within 1 day
Audit logs*		Login attempts, data modifications, configuration changes	Login attempts, data modifications, configuration changes
Support	Direct support on Discord	Dedicated enterprise support	Dedicated enterprise support
Data privacy	Limited data training with the option to opt out	Perplexity and third party LLM partners do not train on your data	Perplexity and third party LLM partners do not train on your data
Certified security	SOC 2 Type II	SOC 2 Type II	SOC 2 Type II

Is it safe to use Perplexity?

Yes—with Perplexity Enterprise Pro for Finance, but caution is advised.
Perplexity offers a specialized solution for financial institutions with enterprise-grade protections, but some concerns remain.

✅ Enterprise-Grade Features

- **SOC 2 Type II compliance** for data security, availability, and confidentiality.
- **Zero Data Retention Policy** for API interactions—no prompts or responses are stored or used for training.
- **Enterprise Pro for Finance** includes:
  - Cited answers from SEC filings, earnings calls, and proprietary databases.
  - Deep research capabilities for due diligence, ESG analysis, and market trends.
  - No training on customer data.

⚠️ Risks and Limitations

- **Consumer-facing versions** may retain data and lack strong authentication protocols.
- **Prompt injection vulnerabilities** have been identified in research environments.
- **No built-in multi-factor authentication (MFA)**, unlike competitors (SSO/SCIM exists, though)

Can I trust Perplexity?

Yes—with enterprise deployments and proper governance, but transparency gaps exist.

🔐 Privacy Practices

- **Enterprise data is never used for training**, including data shared with third-party providers like OpenAI and Anthropic.
- **Uploaded files are retained for 7 days**, with admin controls for deletion and access.
- **Incognito Mode** can be enforced across organizations to disable search history.

🛡️ Security Architecture

- **Vulnerability Disclosure Program (VDP)** encourages responsible reporting of security issues.
- **Trust Center** provides access to compliance documentation and security updates.

⚠️ Concerns Raised by Experts

- **Ambiguity in data usage policies** in earlier versions raised GDPR/CCPA compliance questions.
- **Weak authentication protocols** increase risk of credential-stuffing attacks.
- **Debugging features and developer exploits** were found to be vulnerable in past assessments.

Perplexity Common Use Cases

SBS INSTITUTE

- 🧠 Best-in-class Research Capabilities
- 🌐 Threat Intelligence and Current Events Monitoring
- 📊 Investment Research & Market Analysis
- 📈 Real-Time Stock & Earnings Insights
- 📝 Due Diligence & Deal Evaluation
- 🤝 Strategic Scenario Planning
- 📄 Regulatory & ESG Intelligence

USE PERPLEXITY IF YOU...

- Want to leverage Perplexity Enterprise Pro for Finance and do deep research into specific sets of data at your institution
- Probably better served being a “build” model unless you want to develop some specific internal use cases around data analytics and not pay to host your own AI model

DEEP SEEK

- **DeepSeek is known for being lower-cost and one of the first models to include deep reasoning…**
- **Question**: Is it safe to use DeepSeek if I'm a highly regulated financial institution?
- **Question**: Can I trust DeepSeek to keep my data safe and private?
- **Question**: Why would I choose to use DeepSeek?

deepseek

Access Type	Pricing	tails
Local Model Use		Dow un models locally (e.g., via LM Studio, O atbox).
Model License	ree / ght	Permissive lic g., MIT) allow research and commercia l
Web Chat (Free Tier)	Free	requests/day dit card required.
Web Chat (Pay-as-you-go)	$4.90 for 200 ts	thly c ent, prepaid credits.
Web Chat (Pro Plan)	$ th	Unl sts, priority support, analytics.
API Access	$ Pay-pe	er million tokens, no subscription required.

Is it safe to use DeepSeek?

Is it safe to use DeepSeek?

NOT RECOMMENDED for direct use in regulated environments without significant safeguards. While DeepSeek offers powerful open-source models and cost-effective performance, its **privacy, legal, and infrastructure risks** make it unsuitable for sensitive financial data unless self-hosted with strict controls.

⚠️ Key Risks for Financial Institutions

- **No enterprise-grade version available**: DeepSeek does not currently offer a dedicated enterprise deployment with contractual guarantees or indemnification.
- **Data retention concerns**: DeepSeek's online platform may retain user data indefinitely, even after account deletion.
- **Servers located in China**: All data is stored under Chinese jurisdiction, raising concerns about compliance with U.S., EU, and other international data protection laws.
- **No indemnification**: Users bear full legal risk for outputs, including potential IP infringement.
- **Censorship and bias**: DeepSeek may suppress politically sensitive topics due to Chinese regulatory influence.

✅ Safer Use Case

- **Self-hosted deployments**: Running DeepSeek locally (offline) avoids cloud-based data sharing and offers more control.
- **Open-source flexibility**: Developers can inspect and modify the model but must implement their own security and compliance layers.

Can I trust DeepSeek and Hangzhou?

Not without significant reservations. DeepSeek's privacy policy and operational practices raise multiple red flags for enterprise trust.

🔐 Privacy Policy Highlights

- **Broad data collection**: Includes prompts, chat history, device info, and location.
- **Data used for model improvement**: Inputs and outputs may be reused to improve services unless explicitly opted out.
- **Third-party sharing**: Data may be shared with analytics, advertising partners, and law enforcement.
- **Keystroke tracking**: Privacy policy references monitoring user keystrokes, raising concerns about behavioral profiling.

🛡️ Security Concerns

- **No clear encryption standards**: No mention of AES-256 or TLS in documentation.
- **Weak or absent guardrails**: DeepSeek lacks robust moderation and safety filters, making it vulnerable to adversarial attacks.
- **Unclear training data sourcing**: No transparency on whether training data respects IP or privacy laws.

⚠️ Regulatory and Legal Risks

- **Subject to PRC law**: Disputes must be litigated in China, with limited recourse for foreign entities.
- **Banned by multiple governments**: NASA, U.S. Navy, Taiwan, Italy, and Texas have prohibited DeepSeek use due to privacy concerns.

USE DEEPSEEK IF YOU...

- Are selling data and corporate secrets to China
- Want the Chinese Government to live in your home or corporate network, steal your data, and siphon your intellectual property

OPENCLAW

You may have already heard of OpenClaw – the revolutionary new, FREE, LOCALLY HOSTED AI model that is like an AI Agent on steroids. It w[illegible]r apps already, and comes with pre-packaged “skills.” It’s a programmal[illegible] worker tha[illegible]ransform your AI chats into actions. It’s super cool – but NOT ready fo[illegible] institution prime[illegible]et… for a variety of reasons.

Feature	ChatGPT	OpenClaw
Executes commands		Yes
Access to files	No	Yes
Runs workflows	No	Yes
Multi-step reasoning	[illegible]	Built-in
Works across apps	[illegible]	Yes

AI BOT BATTLE CONCLUSION

WHAT DOES SBS RECOMMEND?

- ✓ Depends a bit on your use-case… and the Build vs. Buy concept
- ✓ BUT – if you are just exploring AI productivity and efficiency gains for employees… AND you want to use member information, consider the following:
  - **MICROSOFT COPILOT is the natural first step – you're already in it**
  - **ChatGPT** is the most popular for a reason, but your data is leaving your environment
  - If you are a Google shop, **Gemini** would be your natural first step… otherwise avoid for institutional use
  - **Perplexity** and **Claude** have specific use-cases, but not a lot of strong pros to take data out of your environment…
  - **Don't touch DeepSeek, please…**

BIG TAKEAWAYS

- Yes – you can use AI!
  - However, you need to be able to justify your use-cases and provide the typical governance documentation expected by regulators…
  - Do your homework (AI can help with that!)
- Consider the Build vs. Buy scenarios
- Don’t use free AI models for work-related tasks
- Premium/Enterprise versions of AI platforms are generally safe to use… depending on your institutional use cases
- Please never use DeepSeek!

DOWNLOADS & MORE!

- HEAD TO OUR LANDING PAGE AND DOWNLOAD SOME GOODIES!
- INCLUDING:
- CHANCE TO WIN A FREE SBS INSTITUTE WEBINAR OR MEMBERSHIP!
- TODAY'S SLIDE DECK(S)!
- PRESENTATION SURVEY - WE LOVE FEEDBACK! TELL US HOW WE DID!
- FREE DOWNLOADS - AI POLICY, BLOG POST, INFOGRAPHIC, MORE!
- SIGN UP FOR IN THE WILD!

HTTPS://SBSCYBER.COM/DAKCU

SBS INSTITUTE – WE'VE GOT YOUR CONTENT

https://learning.sbscyber.com/

FALL 2025
High Performer

SBS - TOP RATED CYBER GRC COMPANY

Clients Love Us

TRAC: CYBER RISK MANAGEMENT SOFTWARE

TRAC
From SBS CyberSecurity

Frustration-Free Risk Management

Software Advice 4.8
GetApp USER REVIEWS
Capterra 4.8

- Vendor
- IT
- Business Continuity Management (BCM)
- Action Tracking
- Audit
- Bank Secrecy Act (BSA)
- Commercial Account Tracking (CATRAC)
- Compliance
- Enterprise Risk Management (ERM)
- Information Security Program (ISP)
- NIST CSF

https://sbscyber.com/

COMPLIMENTARY RESOURCES

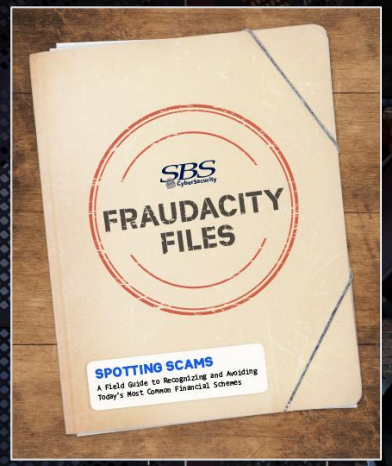

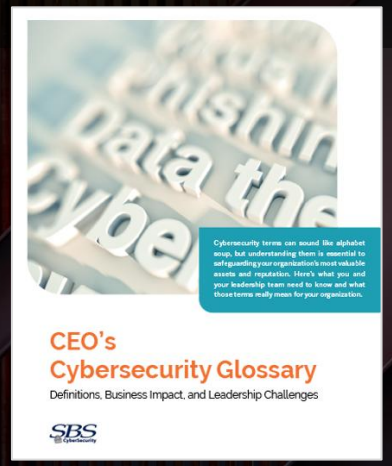

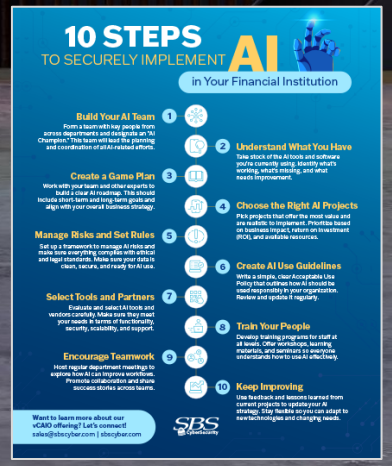

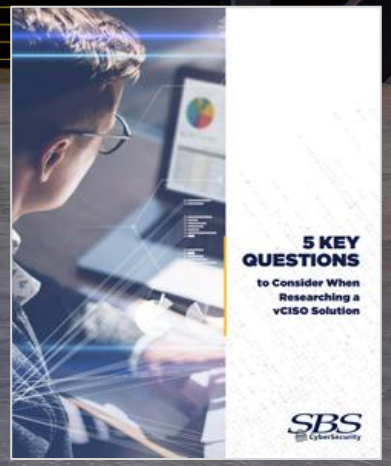

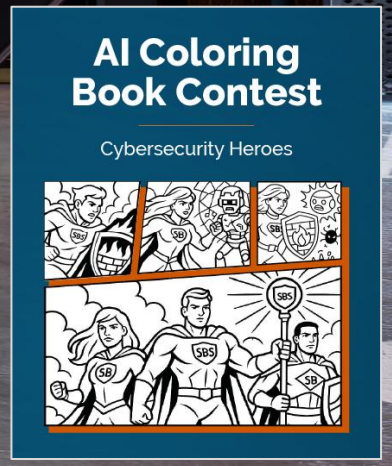

https://learning.sbscyber.com/resourcelibrary

Follow SBS on LinkedIn

Stay in the know with expert tips and insights, upcoming webinars and events, blog updates, and more.

https://www.linkedin.com/company/sbs-cybersecurity

Email Jon to receive our weekly top-secret cybersecurity newsletter… or click here:

SIGN UP TODAY

JON WALDMAN

- President, Partner, Co-Founder – SBS CyberSecurity
- CISA, CRISC, CPDSE
- Masters of Information Assurance, Dakota State University
- Mission: help you make **empowered** cybersecurity decisions
- Phone: 605-380-8897
- jon@sbscyber.com
- www.sbscyber.com

SBS Institute

- sbsinstitute@sbscyber.com
- 605-269-0909

Follow us on Social: