KEEP CALM & PROMPT ON

DeepSeek 101: Security & Risk Management Cheat Sheet

DeepSeek is a Chinese AI company that emerged in December 2023 and is quickly becoming a major competitor in the AI space since **the release of their R1 model** (January 2025). With 671B parameters, it matches OpenAI o1 on many benchmarks and ranks third globally in performance.

What sets them apart is their innovative approach to model development - they claim to have built their model for just $5.6M (compared to typical industry costs of $100M+) by leveraging pure reinforcement learning and optimizing GPU usage.

While they offer significant cost savings with API prices up to 90% lower than competitors, there are some concerns. The platform follows Chinese content regulations, and their terms grant them broad rights over user-submitted content and AI-generated outputs.

Its open-source model is a major win, because it allows companies to deploy AI locally, keeping sensitive data in-house and complying with even the strictest security policies. That's a huge shift, especially when considering the amount of sensitive information people are sharing with AI systems.

DeepSeek is gaining a lot of momentum, but the possibility of backdoors or vulnerabilities in the technology isn't something to take lightly. It's important to note though that these risks aren't unique to DeepSeek. Organizations must take a comprehensive approach to securing data, no matter which model they're using.

Can I test DeepSeek?

Yes, you can try it out at chat.deepseek.com. However, you'll need to create an account first, which differs from some other AI providers that offer immediate access.

How do I block the use of DeepSeek?

- Block the domain deepseek.com and its subdomains using your URL/Domain filtering solution.
- If you are managing employee mobile devices you can use your MDM to block the DeepSeek mobile app.

How do I safely use DeepSeek?

- Monitor the usage of your employees by using a browser extension, an agent or a network inspection solution.
- If you are using the DeepSeek API in your applications, use an AI Gateway to inspect and protect all prompts and responses.
- If your developers are using the DeepSeek models in their code assistants, have an agent or a network solution to redact secrets or PII in prompts and to validate that the source code returned in the responses is not malicious or vulnerable.

How do I coach my employees?

- Send formal notices to employees about AI security measures.
- Implement clear data privacy guidelines to protect against leakage of sensitive data.
- Implement a security solution to show an education popup to employees sharing sensitive information.

Feature Comparison Table

Feature	OpenAI o1	DeepSeek R1
API Cost (per 1m Tokens)	$15.00 (input), $60.00 (output)	$0.55 (input), $2.19 (output)
End User Cost	200$/month	Free
Context Window	200K tokens	128K tokens
Maximum Output Tokens	100K tokens	32K tokens
Compliance & Safety	Strong safety guardrails	Vulnerable to jailbreaking, less robust guardrails
Release Date	December 5, 2024	January 21, 2025
Open Source	No	Yes (MIT License)
Model Parameters	175 billion total	671 billion total, 37 billion active per token (Mixture-of-Experts)
Training Data	Extensive datasets (books, code, etc.)	Trained on 14.8 trillion tokens with RL techniques
MMLU Overall Accuracy	91.8%	90.8%
Customizability	Limited Customization (Black Box)	Full Customization (Open Source)
Deployment	Cloud-Based	Cloud or On-Premise Deployment
Web security	CSP, XFO, Referrer policy	Weak security. No permissions policy.