

Which FortressAI capabilities are generally available, and which are in beta?

Short answer: Core kernel-level enforcement — blocking sensitive data from reaching AI tools, per-tool authorization, and location-based data policy — is generally available. Browser-based AI usage detection and enforcement, Microsoft Purview/MIP sensitivity label enforcement, and automatic PII/PDPL anonymization are currently in beta.

Generally available

- **Kernel-level data protection** — sensitive data is blocked from leaving through ChatGPT, Copilot, Gemini, or any other AI tool, enforced at the operating system's deepest layer.
- **Per-tool assessment** — granular control over exactly which AI tools are authorized and which are blocked.
- **Location-based enforcement** — policy tied to the data locations that matter, so source code stays in the repository rather than in a chatbot.

In beta

- **Browser-based AI usage detection and enforcement** — monitoring and enforcing policy against browser, page, and extension activity.
- **Microsoft Purview / MIP sensitivity label enforcement** — extending existing Purview classification into per-AI-tool policy, so organizations that already classify data don't maintain a second scheme.
- **Automatic PII / PDPL anonymization** — dynamic filtering that strips personally identifiable information from otherwise-allowed prompts. This supports the middle path most organizations want: let people use AI productively, but remove the sensitive elements automatically rather than relying on employees to self-censor.

Beta capabilities are functional and in active use, but should be evaluated in your environment before you depend on them for a control. For current beta availability and enrollment, contact your Cyber Crucible representative.