

Frontier AI and Colorado's Cybersecurity Response

A presentation to the Joint Technology Committee
How OIT is Adapting
June 2026



What changed - and why it matters

A new class of AI – frontier models – can now autonomously discover software vulnerabilities at a scale and speed that previously required expert security teams working for months.

Project Glasswing

Roughly **50** major technology partners provided early access to Anthropic's frontier model to identify vulnerabilities before adversaries develop similar tools.

- Program has already identified more than 10,000 critical or high-severity vulnerabilities across every major operating system and web browser.

Critical vulnerabilities: Flaws an attacker can exploit remotely, with no login required and no action needed from a user.

High vulnerabilities: Nearly as severe as critical but may require an attacker to already have some level of access or await a specific condition.



OIT's Response

Why this a fundamental shift

The vulnerability lifecycle - then vs. now

Pre-2018	2024	2026
Disclosure → Exploit 63 days Human teams hunting for months or years	Disclosure → Exploit 5 days All tools compress the window sharply	Disclosure → Exploit 10-15 minutes Exploit creation from time patch is released

After Mythos, **10,000+ critical/high vulnerabilities queued simultaneously**. Patch avalanche incoming.



The cost barrier
that protected most
organizations is
gone



Discovery is no
longer the
bottleneck



Decades of hidden
vulnerabilities are
now discoverable at
once



The time-to-exploit
window has
collapsed

What we're doing about it

Already in motion


Full scanning of all state systems

30-day remediation of all critical and high vulnerabilities, focusing on internet-facing systems

Testing 24-hour internet-facing response to critical and high vulnerabilities

July 1, 2026

Surge expected from the Project Glasswing disclosure wave



Adversaries are estimated to be roughly six months behind Anthropic in comparable AI capability

Colorado is armed and actioning

01 Endpoint Protection

24/7 managed detection and response across endpoints. CrowdStrike analysts active hunt threats and respond on Colorado's behalf. In March 2026, unveiled Agentic MDR - AI agents that automate high-friction workflows to stop AI-accelerated threats at machine speed.

★ Project Glasswing partner, meaning they receive Mythos vulnerability intel ahead of public disclosure – a strategic advantage.

03 Anthropic POC (Claude Security)

Evaluating Anthropic's AI-powered vulnerability scanning tool (public beta released May 2026). Unlike traditional scanners, Claude Security can reason about how vulnerabilities combine into attack chains – directly addressing the chaining problem our other tools cannot see.

Note: We've requested approval from FEMA/CISA to use SLCGP grant funding for the next phase of Frontier AI use at the State.

02 Vulnerability and Attack Surface Management

Scanning foundation that identifies and prioritizes Critical/High findings. AI-powered Vulnerability Priority Rating that forecasts which vulnerabilities are most likely to be exploited – not just static severity scores.

Tenable One maps directly to CISA BOD 26-04 four-factor risk model.

04 Claude Enterprise (Trial)

Being evaluated for analyst productivity and operational support. Assists analysts in triaging vulnerability reports, drafting compensating control documentation, and accelerating patching workflow communications that an aggressive response posture demands.

Not a vulnerability scanner – reduces the human friction in the remediation loop.



Limitations and gaps

Current stack limitations

Tenable finds known vulnerabilities

Cannot reason about how vulnerabilities combine into attack chains.

CrowdStrike

Only as effective as its sensor deployment.

Anthropic POC / Claude Enterprise

Still in evaluation. The gap between 'evaluating' and 'protecting us at scale' is where the risk lives.

Internet-facing remediation

Cost, vendor engagement timelines, and business impact constrain how fast we can close external exposures.

Coverage gaps

Patch automation

A 24-hour SLA is the right future target but will require automated deployment pipelines. Manual patching will not hold under volume.

Sensor coverage completeness

Older systems, contractor-managed systems, or systems stood up outside standard processes could be blind spots.

Vendor contract remediation

Third-party operators may not be held to the same patching standards as state systems. Some contracts may need to be rewritten for a more aggressive SLA.

Identity & access (non-human)

Machine identities - service accounts, API keys, tokens - now outnumber human users and are the leading cloud breach vector.

Thank you.

Questions?

