AppSecAI

Your Time Is HERE

Session Notes — OWASP Bay Area Chapter Meeting, April 16, 2026

SPEAKER	EVENT	DATE	RESOURCES
Bruce Fram, CEO, AppSecAI	OWASP Bay Area Chapter	April 16, 2026	appsecai.io/owasp-april-2026

Key Points

20 hrs
TIME TO EXPLOIT (WAS 2.3 YEARS)

252
DAYS AVG TO FIX A VULN (VERACODE)

8.2 min
FIX TIME WITH AI-ASSISTED WORKFLOW

The moment has changed

Anthropic's Mythos — a next-generation AI model capable of autonomously discovering thousands of vulnerabilities in production code — put AppSec in the New York Times and triggered an emergency meeting called by the Treasury Secretary and Fed Chair with top bank CEOs

Jamie Dimon, CEO of JPMorgan Chase and the most prominent banker in the country, is talking about application security on Yahoo Finance. The C-suite cares now.

Time-to-exploit has collapsed from years to 20 hours. Average fix time is still 252 days. The math doesn't work.

The new AppSec operating model

Developers don't code anymore — they ask. Five new dev tool announcements in the last 24 hours alone.

The AI model is like the internal combustion engine — the engine alone is interesting, but it's the context you wrap around it that creates value. Building a car is very different

from building a water pump. The same LLM needs very different context for security than for legal or coding.

Shift left has failed. 20 years of evidence. The "tool tornado" on the dev side is not your problem to solve.

AppSec must move from influence to agency — from advising on risk to taking direct action on what goes out the door.

Nobody ever got promoted for fixing a SQL injection. The incentives will never align. Security must own the risk.

Manual and manual-assist won't scale. The path is: Manual → Manual Assist → Exception Auto → Full Auto.

> *"The question is no longer whether AI can find the bugs. It can. The question is whether we can fix them before the attackers get there."*
>
> **— Chris Hughes, CEO of Aquia, founder of Resilient Cyber**

The GRASP Framework

Governance — Who controls the AI? Hard lesson: Uber blew their entire annual token budget in April.

Reliability — Models have outages. Anthropic killed a model version with two days' notice and broke our product.

Assurance — Are results accurate? They'll tell you everything is perfect. Benchmark and verify.

Scalability — Does it work across the enterprise, not just a pilot?

Protection — It's not just prompt injection. Credential management across AI systems is basic, hard, and unsolved.

Your job now

Build the system around the model: orchestration, validation, integration with your tooling and dev process

Set the policy: what's auto-fixed, what needs review, what's blocked

QA the output: validate AI decisions, catch drift, ensure fixes are correct

Measure the outcome: backlog shrinking? Developers adopting? Posture improving? Prove it.

Skills for the new world

Curiosity is the most important skill. You came out on a Thursday night — you have it.

Traditional AppSec skills still needed, plus: AI/ML literacy, policy design, building tools, metrics, solution evaluation, AI output QA

You can't evaluate one new technology per quarter anymore. Five new things come out every day.

Three moves for Friday morning

Sign up for Project OASIS — Review AI-generated security fixes for open source. 15 min/week. Your name on vetted PRs.

Try the GRASP skill — Run it against any AI tool your team uses. Automated scoring.

Download and read the book — The economics, the broken process, and how to manage up.

Full Transcript

> **Slide 1 — Your Time Is HERE**

Thank you. Hey, thanks, Sai. Good evening.

I'm Bruce Fram, CEO of AppSec AI. I've been CEO of seven different enterprise software companies over many years. Just can't seem to stop doing it. Some of you may know me or know Contrast Security. I was the founding CEO of Contrast, which at one point was a billion dollar unicorn. And I'm sure many of you have used the Contrast products.

At Contrast, we were very good at finding vulnerabilities. I've come back to help fix them. That's what AppSec AI does — we fix vulnerabilities at scale. Honestly, before AI, code remediation was a dead end. It just wasn't going to be effective. But right now, our time is here in application security because of everything that's happened with AI.

This talk is not theoretical. I hope all of you will take action after this. I'm going to talk about a number of different areas — ways to look at the world. It is absolutely about, hey, here's what I can do tomorrow morning. And that's what we'll talk about today. And I hope we have some fun doing this because, boy, things are changing.

Slide 2 — Hi Mom & Dad — I am an AppSec Engineer!

The first thing is — hey, mom and dad, I'm an AppSec engineer. No one knows what we do. In the end, AppSec is a great business. It's super important. But let's not kid ourselves. No one knows what it does.

These are headlines from last Sunday's New York Times. They're talking about application security in the New York Times on the Sunday paper — and where I come from in New York, that is the epitome of getting to the top of the world. This is largely because of what Anthropic announced last week with Mythos. Can you just raise your hand if you know what Mythos is? Okay, about two thirds of you do.

This has created a huge spotlight on AppSec. If you don't know what Mythos is — Anthropic has announced a future model that has been incredibly effective at finding thousands of vulnerabilities in code. They've released the model. They've put together an entire project with open source and corporate developers to go out and fix all these vulnerabilities before they can be exploited — not only by Anthropic's future models, but lots of others. Anthropic might be ahead now, but I think we all know this is a fast moving space.

So the crazy thing is your mom or your dad or your partner who don't know anything about AppSec — you're now in the New York Times.

Slide 3 — The CEO/Board REALLY Cares...

The good news for us in AppSec is the CEO and the board really care. In the past, let's not kid ourselves, they didn't really care. The CISO would go before the board of directors meeting once or twice a year, show a bunch of charts and graphs and bugs and KPIs. I'm sure their eyes glazed over. But with the Mythos breakout, they really care.

Here's some evidence. On April 7th — barely a little more than a week ago — right after this was announced, the Treasury Secretary and the Fed chair called a meeting of all the top banks in New York City and had them attend personally. They're talking about application security. That is crazy.

If you see here on the bottom, we've got Jamie Dimon, the CEO of JPMorgan Chase, talking about Mythos. For those of you not in the financial sector, Jamie Dimon is the most prominent banker in the country. He steered JPMorgan Chase through the 2008 crash. He's the longest standing CEO of any big New York bank, incredibly well regarded. And he's talking about Mythos. This guy is talking about application security on Yahoo Finance. I never thought that day would ever come. I'm sure you didn't either. But they're there.

Slides 4–5 — Because they Understand This / Can Collapse the Economy

Why do they care? Because they appreciate this — from the Zero Day Attack website, which just came up about a month ago. It wasn't long ago that a vulnerability was discovered and it would be a couple of years or a year before it was exploited. That has collapsed to 20 hours. That doesn't mean every vulnerability is going to be exploited in 20 hours, but it's certainly not going to be a year. This chart is based on real data.

Software runs everything. And software has vulnerabilities, always has vulnerabilities, and we're not fixing them fast enough. Time to exploit was 2.3 years. Now it's 20 hours. The average time to fix a vulnerability is 252 days. This data is from Veracode. If you're a Veracode customer, you actually care. And it's taking you nine months to fix things. That's why they called the meeting.

Slide 6 — Great for US. Might Even Be Fun!

This is great for us — the application security industry — and it might even be fun. I started the company with my co-founder, Michael Cartsonis, who also helped start Contrast, because this is fun. We're having fun. I hope you have fun.

The code explosion has created very strong demand for what we do. Whether it's first party code, libraries, infrastructure as code — this is what it's all about. And we're very well positioned.

A CISO I know said, "Hey Bruce, do you think I should tell my younger colleagues and my kids that security is a good place to go?" I said, "No, man. Security is a great place to be if you've got the right skills."

Slides 7–8 — How do we deal with this? / Do you feel confused?

The real question is — how do we deal with this and win? I'm going to cover: (1) this is not the first technology transformation, (2) a new DevSecOps process is coming, (3) what does automation mean in our business, (4) a framework to think about our jobs, and (5) actions you can take — including a new open source project called Oasis launching today.

Do you feel confused, unsure, behind? I do. Let me be very clear. Whether you're here or working in one of the big AI labs, nobody can keep up. New things are coming out every day. You just can't keep up. That is the bad news.

The good news is you're not behind. This is about figuring things out together. Nobody has the answers. I don't have the answers. But there is a lot we can do together. And that's what it's about.

Slides 9–10 — We've Seen This Before / Same Story, Different Technology

The internal combustion engine, invented in 1876, spawned aircraft, manufacturing, automobiles, logistics, trains. No one saw that when they invented it. It took 30 years to become common. It changed all kinds of jobs. It changed everything in the economy.

Today, we've got the same story with a different technology. AI and LLMs are going to change every industry. It's not just coding — it's drug discovery, vehicles, supply chain, legal. But the value we add is what we put around it. It's the context you put around it. This is like an engine — and how you use it, the context and knowledge you put around it, that's the value.

Anthropic comes out with announcements on cybersecurity. Cybersecurity stocks drop 10%. They're not going to put Zscaler out of business. All of us know that. But the market doesn't know that because it just doesn't have enough context around it.

Slides 11–12 — Devs don't code anymore / The more they automate

Developers don't code anymore. They ask. And they use a lot of different tools to do this. It changes all the time. Nico from Alien Giraffe is not coding in a classic way. This has massive implications for AppSec.

In the last 24 hours, there's been five new announcements that devs have to think about: Google Gemini Mac app, OpenAI Codex integration with Claude Code, new Codex desktop that reads your screen, new Claude Code, and Anthropic announcing Opus 4.7. This is what's happening to developers every day.

The more developers use AI, the more they automate. Anthropic's data shows developers auto-approve 20% at first, rising to 40-50% after a thousand sessions. This curve has huge implications for application security.

Slides 13–14 — AppSec is manual and slow / We never catch up

Application security is a manual process and slow. None of these processes are free. This is an 11-step process that it takes to fix something in code inside an enterprise. Every step takes time, has a cost, and has people. It doesn't meet the moment.

With this manual process, we never catch up. If you've got a vulnerability backlog that is growing, you never clear it. If you have a 50% clear rate — which is incredible, and nobody has got that — you're done in five years. And I'm talking about mediums and highs. People have thousands and tens of thousands of these.

Slides 15–16 — The Vulnpocalypse is coming / Nobody fixes them

With Mythos, something called the vulnpocalypse is coming. Finding is easy. Our theory is finding vulnerabilities is easy. We don't need more. It's going to be free. Chris Hughes says, and we agree 100% — "The question is no longer whether AI can find bugs. It can. The question is whether we can fix them before the attackers get there." And the answer is we're not.

It costs an enterprise five to $20,000 to fix a medium or high vulnerability because there's 11 steps. Now we can get that down with AI to 1/100 of that cost. We were just doing a customer the other day and they had an 85% false positive rate. You can get that way down. We have to move AppSec from the department of no to an acceleration team.

Slides 17–18 — Direct action / Someone has to own what goes out the door

The new model — you're going to take direct action. Today, you're being blocked because the devs own the code, the business does not want to invest in security, and tools give you 40-80% false positives. The situation we've been in is: you have influence, but not agency. We've got 20 years of failure with that.

We have to get the incentives aligned. Developers get rewarded for shipping features, function, and performance. Nobody ever got promoted for fixing a SQL injection. Security is risk. Managing risk is going to become our issue.

Someone has to own what goes out the door. That's going to be us — but us with automation. Your job is to figure out the path from manual to exception auto to full auto... and put in place the policies. Manual and manual assist won't work anymore.

I had a big debate with my co-founder Michael Cartsonis about full auto. I said "Michael, don't do that. People aren't ready." Michael was right and I was wrong — because full auto is coming, or at least exception auto. The 20th SQL injection you've seen, you don't need to look at that every time.

Slides 19–20 — Is this fast enough? / Your Job

Pre-Mythos, manual assist: we can get it down from 252 days to 8.2 minutes. That is really cool. Our customers love it. But is that fast enough? I don't know. We're going to find out.

Your job is putting in place the systems around the model. The real value isn't the frontier model itself. It's the system you build around it — orchestration, validation, integration with your tooling, your dev process, your security expertise. Building the system, setting the policy, QA-ing the output, measuring the outcome.

Slides 21–23 — The GRASP Framework / Categories / Matrix

We learned a lot of hard lessons and put in place something called GRASP — Governance, Reliability, Assurance, Scalability, Protection.

Governance: who's governing the cost? Uber spent their entire annual token budget in April. Kevin, our CTO, blew a couple thousand bucks in tokens with no governance.

Reliability: models have outages. Anthropic killed the Haiku model version we used, with two days' notice, and our product broke. And results: we don't move to a new model every time something comes out because results will change. You can't just roll that out to your enterprise.

Assurance: are the results accurate and relevant? They'll tell you everything's perfect. Believe me. Benchmark and verify.

Scalability: can it handle workload, multiple time zones, across your organization? We've all had lots of issues.

Protection: beyond prompt injection. Nico from AlienVault talked about credentials. With all these AI systems, credential management is basic, hard, and unsolved.

Slides 24–25 — Enterprise fit / Applied: OpenClaw

You have to implement this framework across layers: the LLM core, the runtime/agent layer, the domain application, and your enterprise context. No LLM is going to solve your enterprise context — is this coded the way we do it? Our standards? Our business rules? Our compliance? Putting that together is super important.

For fun, I created a GRASP skill and ran it against open source projects. You can see the results on OpenClaw — pretty accurate. It evaluates the repo and asks you three to five questions, not 100. You can get a really quick view. Obviously, OpenClaw didn't score too well on it.

Slide 26 — Skills You Need for the New World

The biggest skill you need is curiosity. You came out on a Thursday night — I appreciate this — to learn about what's going on in AI and application security. That is the most important thing.

We still need traditional AppSec skills, but we need new ones — understanding what models do and how they work, policy design for automation, building tools with AI, metrics, solution evaluation, and AI output QA.

In the old days — "we're going to look at this new technology next quarter." That's not going to cut it with five new things coming out every day. You have to quickly separate the hype from the substance. Is it material or the flavor of the day? You can't just look at one thing a quarter anymore.

Slides 27–28 — Tropicon / Project OASIS

So you may wonder why I'm showing you a picture of a Mexican resort. My co-founder, Michael Cartsonis, is in Cozumel at Tropicon rolling out an open source effort that you can all participate in.

We have partnered with Chris Holt from Intigriti, who used to run Intel's bug bounty program. He's created the Open Automated Security Initiative for Software — OASIS. The idea is to fix bugs in open source using both AI and people. Open source project folks get a lot of AI slop. This project is a process to review AI-generated security fixes, validate them, and submit real vetted fixes.

You don't have to know any products. If you can use GitHub, you can participate. 15 minutes a week, one a week. You get your name on PRs. If we get the OWASP community behind this, we can vet hundreds and thousands of these with AI plus your expertise. I think it'll be fun.

Slides 29–31 — The Playbook / Three Moves / Resources

I've got my book here. Post-Mythos, it applies. This is not a book about products or technology. It's about the process and the costs. It's a quick read with lots of checklists.

Friday morning — three moves: Sign up for Oasis. Try the GRASP skill on a project and give me feedback. Download and read the book. All free. Everything is at one scan.

Q&A — Audience Questions

Q: Can't developers do this? Won't AI fix this stuff and shift left?

This is not going to be a popular opinion, but shift left has failed. It's been 20 years. It has failed. Not all code is being developed in-house. Developers are focused on every new tool that comes out — which creates other vulnerabilities. The "tool tornado" is real. Every study that's looked at AI-generated code shows it generates a ton of vulnerable code. Maybe AI will get very good at fixing simple SQL injections at the root. But it will not fix things within your enterprise context. We're a long way off from that.

Q: Doesn't Dev own the risk?

In the end, we own the risk. Shouldn't Dev own some of that? Yeah, but they're never going to own it. We have to stand up and own the risk. The developer is not going to look at a two-line SQL injection fix. We're going to have to own all that stuff. Of course, there are exceptions for more critical things where you'll jointly do it. But developers are focused on keeping their jobs — meaning learning all these new tools. We have to guide them. We have to step up and own it and articulate it. The good news is we now have a platform to do that.

Closing

Thank you very much. Go get the book and let's take some action.