

# Jason Colapietro: expert and speaker kit

Jason Colapietro helps teams define useful AI workflows, inspect the sources behind search answers and build factual acceptance checks they can rerun.

## Bios

### 25 words

Jason Colapietro founded Suede AI. He builds agent platforms and publishes reproducible examples of MCP connections, source SEO checks and source-grounded AI answer evaluation methods.

### 75 words

Jason Colapietro is the founder and CEO of Suede AI, where he builds and operates AI agent platforms. His working examples connect an MCP client to Suede Creator Skills, check a published evidence page and evaluate source-grounded answer fixtures. His issued training covers Claude API development, MCP, search, prompting and generative AI. He helps builders and business owners define useful scopes, keep dated source evidence and turn acceptance criteria into checks their teams can rerun.

### 150 words

Jason Colapietro is the founder and CEO of Suede AI, an agent platform company in West Palm Beach, Florida. He builds and operates AI agent platforms and leads Suede AI SEO, the company's search visibility practice. His working examples give technical and business audiences something concrete to inspect: an MCP client that discovers tools and retrieves a workflow, a source-page checker that detects controlled SEO defects, and an answer evaluator that checks authored fixtures against a factual reference. His issued personal training includes Claude Academy / Anthropic courses, Semrush Academy certificates, HubSpot Academy AEO Fundamentals certification, IBM SkillsBuild, Google Skills and Databricks Academy records, plus Microsoft Learn module and learning-path achievements. Jason's sessions focus on choosing a bounded engagement, naming the person who owns the result and recording the evidence behind an acceptance decision. Audiences leave with a practical method they can apply to a workflow, website or AI answer.

[source: Suede AI qualifications, recorded demonstrations and issued personal credential inventory, 2026-10-21]

## Session and article pitches

All agendas are proposed. Confirm format, duration and session terms with the host.

## Give an agent a tool contract your team can inspect

Audience: Platform engineers, AI builders and technical leaders choosing a first tool-connected workflow.

Formats: Technical talk, Developer article, Hands-on workshop. Recommended: Technical talk.

A useful agent connection starts with a defined task and an inspectable contract. Jason walks through the Suede Creator Skills MCP client: discover tools, search for a workflow, retrieve instructions and handle an unknown skill. The session then translates that example into a pilot specification with input boundaries, permissions, failure behavior and a named operating owner. Participants leave with a checklist they can use to review their next integration.

Takeaways:

- Write a tool contract that names inputs, outputs, authorization and error behavior.
- Retain protocol and source-version evidence so a connection can be reviewed and rerun.
- Define the pilot owner, human approval points and acceptance checks before implementation.

Proposed 30-minute agenda:

- 0–5 minutes: Choose the task and define a useful completion condition.
- 5–15 minutes: Run discovery, search, workflow retrieval and unknown-skill handling.
- 15–23 minutes: Review permissions, source pinning and acceptance checks.
- 23–30 minutes: Apply the checklist to an audience example and take questions.

Demonstration: [Suede Creator Skills MCP client](#)

The recorded MCP client made five protocol requests, discovered nine tools, retrieved a named workflow and verified the unknown-skill error. [source: MCP demonstration results, 2026-10-21]

- Open the pinned server source and identify the declared tool contracts.
- Run the published client and inspect discovery, retrieval and unknown-skill results.
- Compare the report's source commit and returned instruction hash with the recorded example.

Venue fallback: Use the dated downloadable JSON and source side by side if the venue cannot run Node or retrieve the public repository.

Sources:

- [Reproducible MCP, SEO and evaluation demonstrations](#)
- [Recorded demonstration results](#)
- [Pinned Suede Creator Skills MCP server](#)
- [Runnable training-proof examples](#)
- [Issued personal training records](#)

## Turn a source-page SEO finding into a repair you can verify

Audience: Marketing owners, technical SEO practitioners and website teams assigning source repairs.

Formats: Practical talk, Contributed article, Team workshop. Recommended: Team workshop.

Search teams need evidence that leads to an assignable repair. This session uses the Suede AI evidence-page checker to inspect committed page source, introduce controlled defects and verify the same checks after restoration. Jason shows how to connect each observation to a source, owner and acceptance criterion. The workshop closes by separating source-page verification from the dated search and AI-answer samples a team uses to measure later behavior.

Takeaways:

- Record the page address, capture date and complete-response hash beside a source finding.
- Turn title, canonical and indexing defects into explicit repair acceptance checks.
- Track source correctness, answer citations and qualified inquiries as separate measures.

Proposed 45-minute agenda:

- 0–8 minutes: Choose a source page and the buyer claim it supports.
- 8–23 minutes: Inspect the recorded page checks and controlled defect copy.
- 23–35 minutes: Write a repair entry with evidence, owner and recheck.
- 35–45 minutes: Plan a comparable follow-up sample and discuss questions.

Demonstration: [Evidence-page check and controlled SEO defects](#)

The recorded check of committed evidence-page HTML passed eight of eight checks. A controlled copy introduced three defects, and the checker detected each one. [source: technical SEO demonstration results, 2026-10-21]

- Open the source-page report and identify its capture metadata.
- Run the example and inspect the controlled title, canonical and noindex failures.
- Show the restored source returning the recorded checks to pass and draft a repair acceptance entry.

Venue fallback: Present the saved checks and controlled fixture report, with the checker source visible for inspection.

Sources:

- [AI Search Visibility Audit scope](#)
- [Reproducible MCP, SEO and evaluation demonstrations](#)
- [Recorded demonstration results](#)
- [Runnable training-proof examples](#)
- [Issued personal training records](#)

## Build an answer rubric that catches factual drift

Audience: AI product owners, engineering teams and subject experts approving factual answers.

Formats: Engineering talk, Method article, Product-team workshop. Recommended: Product-team workshop.

An answer can read well while changing the facts a buyer needs. Jason uses the published credential inventory as a reference and demonstrates an evaluator that checks authored answers for totals, issuer, award type, source link and relationship. Participants inspect failures that combine course and training achievement totals, remove a verification source or rewrite personal training as partnership. The session ends with a rubric, case-record template and review workflow for their own factual answer type.

Takeaways:

- Turn approved reference facts into named checks with accept/review rules.
- Use controlled failures to test an evaluator before interpreting a model run.
- Keep prompts, settings, source versions and raw output with real model evaluations.

Proposed 45-minute agenda:

- 0–8 minutes: Choose the business decision and approved source facts.
- 8–23 minutes: Run the authored fixtures and inspect expected decisions.
- 23–35 minutes: Draft a case, severity rule and human review condition.
- 35–45 minutes: Define a real model-run record and discuss questions.

Demonstration: [Source-grounded answer fixtures and evaluator](#)

Four authored answer fixtures produced four expected decisions against eight checks per fixture. Failures covered conflated award totals, a missing source and training rewritten as partnership. [source: authored answer-fixture demonstration results, 2026-10-21]

- Read the source prompt and credential inventory.
- Run the authored fixtures and trace the failed checks to the altered fields.
- Open the report, distinguish fixture validation from a model run and draft the metadata needed for a real evaluation.

Venue fallback: Use the saved fixture JSON and report; ask the audience to identify the changed fact before revealing the failed check.

Sources:

- [Reproducible MCP, SEO and evaluation demonstrations](#)
- [Recorded demonstration results](#)
- [Runnable training-proof examples](#)
- [Issued personal training records](#)

## Demonstration setup

Use Node 22 or newer, an approved demonstration environment and the pinned public skill-pack source. The runner executes the three recorded example types. The published result is dated; rerun the current repository before presenting and retain the new report. Add --live only when a current evidence-page comparison is part of the demonstration.

```
git clone https://github.com/JasonColapietro/suede-seo.git suede-seo-demo
git clone https://github.com/JasonColapietro/suede-creator-skills.git suede-creator-skills-demo
git -C suede-creator-skills-demo checkout --detach 46445415a767df5827f5e124f41749ef369cca83
cd suede-seo-demo
node examples/training-proof/run.mjs --skills-repo ../suede-creator-skills-demo --out
/tmp/suede-training-proof
```

[Download the dated report](#) · [Inspect the authored answer fixtures](#) · [Read the source-grounded prompt](#)

## Host introduction

Please welcome Jason Colapietro, founder and CEO of Suede AI. Jason builds and operates AI agent platforms and publishes runnable examples of MCP integration, source SEO checks and factual answer evaluation. Today he will show the method, the source code and the recorded result, then help us apply the acceptance checks to a practical team decision.

[source: Suede AI qualifications and demonstrations, 2026-10-21]

## Interview questions and talking points

### How do you choose a first AI workflow?

Start with a recurring task, representative inputs, a clear completion condition and the person who owns the result. Map the handoffs and exceptions, then specify a bounded pilot with acceptance checks.

- [AI Workflow Assessment scope](#)

### What should a buyer inspect in an MCP demonstration?

Inspect discovery, tool schemas, permission boundaries, result handlers and errors. The published protocol example pairs the pinned server source and instruction hash with the client exchange.

- [Reproducible MCP, SEO and evaluation demonstrations](#)
- [Recorded demonstration results](#)
- [Pinned Suede Creator Skills MCP server](#)

### How do you make a search review actionable?

Attach a finding to a saved observation, an approved source and a repair someone owns. Define the recheck before implementing the fix. The source-page example detects controlled title, canonical and indexing defects.

- [AI Search Visibility Audit scope](#)
- [Reproducible MCP, SEO and evaluation demonstrations](#)
- [Recorded demonstration results](#)

## How should a team measure AI search work?

Keep source correctness, indexing, mentions, citations, answer accuracy and qualified inquiries in separate records. Agree the question set and capture conditions before collection so a follow-up comparison has a defined basis.

- [AI Search Visibility Audit scope](#)
- [Reproducible MCP, SEO and evaluation demonstrations](#)

## How do you evaluate whether an answer preserved a fact?

Approve the source facts and rubric first. Use controlled fixtures to verify the checks, then retain the full prompt, system/model version, settings, date and output when testing a real model. A subject owner resolves the judgments defined in the scope.

- [Reproducible MCP, SEO and evaluation demonstrations](#)
- [Recorded demonstration results](#)

## How does training relate to the work you show?

The issued record keeps course credentials, Microsoft Learn module badges and learning path trophies and GEAR membership separate. Pair relevant personal training with runnable examples and inspectable results. Use separate evidence for company relationships.

- [Issued personal training records](#)
- [Suede AI qualifications and working proof](#)
- [Reproducible MCP, SEO and evaluation demonstrations](#)

## Booking and contact

Share your audience, topic, format, proposed date and desired takeaways. Use the booking link to discuss the session, demonstration setup and any client-specific speaker terms.

[Discuss a session with Jason](#) · [Email Suede AI](#) · [+1 \(561\) 401-0518](#)

[Qualifications and issued records](#) · [Founder and shipped work](#)

Organizer checklist:

- Choose the pitch and audience level.
- Confirm session length, questions, recording permission and accessibility requirements.
- Agree a demonstration environment or use the recorded-results fallback.
- Confirm the current bio, title, source links, contact and any speaker fee in the client-specific booking agreement.