

MASTERS OF SEARCH 2026 • BERLIN

Build, Rank, Convert


Claude Code, from an enterprise web team to a solo side project

Margus Veeber

Head of Web at Pipedrive • Founder at VeeberMedia

pipedrive

>veebermedia



October 22, 2026

I run this at **both** ends

pipedrive

PIPEDRIVE • CONVERSION IS THE JOB

- Every change is a test, not an opinion
- 24 languages, so one change ships 24 times
- Design system and brand are not negotiable
- Legal, privacy and accessibility sign-off
- So: no vibe coding the site

>veebermedia

MY OWN PROJECTS • NO CONSTRAINTS

- Four sites off Webflow to Next.js
- Nobody to ask for permission
- Nothing to break that I cannot fix
- So: vibe coding, hell yes

Same tool . Different **guardrails**.

Everything today depends on which end you are standing at.



PRESENTED BY



First, who is in the room?



PRESENTED BY



WHERE WE ARE GOING

Three parts

01

Vibe coding a site

How I build my own sites, and what usually breaks.

02

SEO and AEO

Schema, FAQs, author pages.
The structural stuff.

03

Conversion

Traffic is down, so this part matters more now.

BEFORE ANYTHING ELSE

Where you actually **run** it

YOUR OPTIONS

- The Claude app
- VS Code
- Cursor
- Conductor
- Straight in the terminal
- And plenty of others

WHERE I RUN IT

- VS Code and Cursor
- Because I want to see the folders and the files
- That is a preference, not a recommendation

Nobody has this figured out. There is no right way to do this.

Not me, not the people posting their perfect setup online.



PRESENTED BY



SECTION 01

Vibe coding a site



PRESENTED BY




WHAT I DO WHEN THERE ARE NO GUARDRAILS

The rules I still keep

- 01 80 percent planning, 20 percent building
- 02 Design one page first. Never the whole site
- 03 Variations in HTML or an Artifact before committing to anything
- 04 Copy rules written down, including what it may not write
- 05 Do not force the design system early. It kills exploration
- 06 Screenshots pasted back into the chat for anything visual
- 07 One page to 90 percent before starting the next

Five versions before **one** build



Exploring variations of one component on the homepage.

WHAT ACTUALLY GOES WRONG

Where it **bites** you

- 01 Lock the design too early and it gets stuck. Nothing better comes out
- 02 It forgets your rules. Design, copy, tone, all of it
- 03 It is confident when it should not be, and it produces bloat
- 04 So: make it assess its own work, every single time

WHAT FAILS SILENTLY

What I always check


- 01 Sitemap – is it generated, and does it update itself when a page is added?
- 02 OG images – is there a default, and does every page get its own?
- 03 Internal linking – is there a system for it written into the rules?
- 04 Tech SEO – does an automated check run before a new page goes live?
- 05 Vercel usage – what runs on every request, and what could just be static

WHY THE RULES MATTER

None of this throws an error when it is wrong. If it is not written into the project rules, nothing catches it. So put the time into the things that have to be followed every single time.

FOUR SITES, OFF WEBFLOW

Why the rebuild was **worth** it



It tanked during the move. Then it came back higher than before.

WHAT GOT EASIER

- Updates take minutes
- New content without fighting a CMS
- Internal linking is easy to handle and improve
- Keeping existing content up to date is easy
- Customer tools drop straight in
- Accessibility issues actually get fixed
- Technical SEO I could never reach before
- Improvements land across every page at once

SECTION 02

SEO and AEO

WHERE IT IS GENUINELY GOOD

The structured layer

- 01 Schema, done the same way every single time
- 02 FAQs from sales calls, reviews and support tickets. Real questions
- 03 Key takeaway blocks
- 04 Titles reformatted to match how LLMs read them
- 05 Author pages, quick to spin up with the schema already on them
- 06 Meta, canonicals, hreflang, sitemaps
- 07 llms.txt

THE PART I ACTUALLY LIKE

You can try something out fast and find out whether it is worth doing, before you put real resources behind it.

ONCE YOU CONNECT IT TO YOUR STACK

What MCPs change

- 01 Connect the CMS – it reads and writes, instead of only suggesting
- 02 Connect analytics and search data – decisions come from numbers
- 03 Changes land across the whole CMS, not one page at a time
- 04 Build a publishing Skill: content in, into the CMS with every field filled

WORTH SETTING UP ONCE

Build Skills for the things you do over and over, and wire the MCPs into the Skill itself. Then it knows where to pull the right data from without being told every time.

Every article, same parts

KEY TAKEAWAYS

KURZ ERKLÄRT

Innerhalb von drei Monaten haben drei sehr unterschiedliche Lager dieselbe Idee auf den Tisch gelegt: Die Website ist kein Projekt mehr, das man launcht und dann liegen lässt, sondern ein System, an dem permanent gearbeitet wird, und die Arbeit übernehmen zunehmend Agenten. Webflow baut dafür die eigene Plattform um (Source, Agent Presence, MCP 2.1). Ploy, gegründet vom Webflow-Mitgründer Bryant Chou, startet gleich agentenzentriert. Und wer mit Claude Code selbst baut, hat diese Schleife heute schon im Repository. Die Frage ist nicht, welches Lager gewinnt, sondern welcher Weg zu welchem Team passt.

☰ DAS WICHTIGSTE

- ✓ Webflow hat auf der Webflow Conf am 2. September 2026 Source vorgestellt: eine Plattform, in der Marketer und KI-Agenten direkt am Code arbeiten. Aktuell nur als Research Preview
- ✓ Ploy ging im Juni 2026 mit 27 Mio. USD Seed-Finanzierung an den Start. Bestehende Website importieren, neue Seiten im eigenen Design, AEO-Vorschläge, ABM-Seiten, alles agentengetrieben
- ✓ Mit Claude Code gebaute Seiten (wie saaswelt.de) sind komplett KI-bearbeitbar, weil alles Code und Markdown ist. Der Preis: kein Dashboard für nicht-technische Kollegen
- ✓ Alle drei laufen auf dieselbe Schleife hinaus: Daten, Einschätzung, Aktionsplan, Mensch entscheidet, Umsetzung
- ✓ Für DACH-Teams zählt Hosting: Webflow und Ploy speichern primär in den USA, nur der Selbstbau bietet EU-Region ohne Umweg

FAQ BLOCK

FAQ

Häufige Fragen zu UTM-Parametern

Was sind UTM-Parameter und wie funktionieren sie? ×

UTM-Parameter sind Textzusätze in URLs, die deinem Analyse-Tool (GA4, Plausible, Matomo) mitteilen, woher ein Besucher kommt. UTM steht für "Urchin Tracking Module" — benannt nach der 2005 von Google übernommenen Analytics-Firma. Heute sind UTMs der De-facto-Standard für Marketing-Attribution. Sechs Parameter sind unterstützt: utm_source (Quelle), utm_medium (Kanaltyp), utm_campaign (Kampagne), utm_term (Keyword), utm_content (Variante), utm_id (Kampagnen-ID, GA4-spezifisch).

utm_source vs. utm_medium — wann nehme ich was? +

Welche Naming-Konvention empfiehlt sich für DACH-Teams? +

Was sind Date-Tokens und wann nutze ich sie? +

Wofür ist der Bulk-Mode gedacht? +

Sind die gespeicherten URLs in der Historie wirklich nur lokal? +

Written once as a rule, then every article gets it.

This used to be the **annoying** part.

Fiddly, repetitive, easy to push to next sprint. Now it is the part AI search reads first.



PRESENTED BY



ONE OF THE BIGGEST BENEFITS

Custom tools and resources

UTM GENERATOR

GENERATOR

UTM-Parameter Generator

Tracking-URLs für GA4, Plausible, Matomo & Co. — mit Historie, Bulk-Mode und Date-Tokens.

Einzel-URL

Bulk-Mode

🕒 Historie (0)

CHANNEL-PRESET

Google Ads

Meta Ads (FB/IG)

LinkedIn Ads

LinkedIn Organic

X (Twitter)

TikTok Ads

YouTube

Newsletter

Podcast

Affiliate / Partner

QR / Print

Gastbeitrag / PR

BASIS-URL *

deine-website.de/landingpage

UTM_SOURCE *

UTM_MEDIUM *

UTM_CAMPAIGN *

UTM_ID OPTIONAL

Tokens: {yyyy}, {yyyy-mm}, {quarter}, {week}

Für GA4-Kampagnen-Reports nützlich

google, newsletter...

cpc, email, social...

produkt_launch_{quarter}

ga4-campaign-id

CONVERSION CALCULATOR

RESSOURCEN

Trial & Freemium Conversion Rechner

Von Signups zu MRR: Übersetze deine Trial- oder Freemium-Conversion in neuen Umsatz - mit Modell-Benchmarks und Steady-State-Projektion.

MODELL

TRIAL OPT-IN

TRIAL OPT-OUT

FREEMIUM

Mehr Signups, niedrigere Conversion, Marktschnitt ~18 %.

TRIAL-STARTS / MONAT	CONVERSION ZU ZAHLEND (%)
1000	10
Top of Funnel pro Monat	Anteil, der zu zahlenden Kunden wird
ARPA (€ / MONAT)	MONATLICHE CHURN-RATE (%)
58	3
Umsatz pro zahlendem Kunden / Monat	Der zahlenden Kunden - für Steady-State

CONVERSION IN MRR BERECHNEN →

ERGEBNIS

Trial Opt-In

NEUER MRR PRO MONAT

5.800 €

100 zahlende Neukunden / Monat · 69.600 € ARR-Beitrag

Calculators, checklists, scorecards. Room to be creative again, and quick enough that you actually try things.

Or a whole directory

LAND

Alle (80)

Deutschland (57)

Österreich (12)

Schweiz (11)

KATEGORIE

Alle

AEO / KI-Analytics

Ausgabenmanagement

B2B-Payments

Backup / Cybersecurity

Banking-Software

Bau-Software

Brand Management

Buchhaltung

Unternehmen

80 angezeigt

Alle Stages

Public

Growth

Series C

Series B

Series A

Seed

Bootstrapped

Acquired



SAP

Weltmarktführer für Enterprise-Resource-Planning

PUBLIC

Enterprise-Software

 Walldorf

Gegründet 1972

↗



TeamViewer

Remote-Zugriff und Fernwartung für Teams weltweit

PUBLIC

Remote Access

 Göppingen

Gegründet 2005

↗



Celonis

Geschäftsprozesse mit KI analysieren und optimieren

GROWTH

Process Mining

 München

Gegründet 2011

↗



Personio

HR-Software für europäische KMUs und Mittelstand

GROWTH

HR-Software

 München

Gegründet 2015

↗



DeepL

KI-Übersetzungen auf Enterprise-Niveau

GROWTH

KI-Übersetzung

 Köln

Gegründet 2009

↗



commercetools

Composable Commerce Plattform für Enterprise-Kunden

SERIES C

E-Commerce

 München

Gegründet 2006

↗

80 companies, filtered by country, category and funding stage. This used to be a project.

saaswelt

SECTION 03

Conversion

You cannot always control the traffic. But you can control the **conversion.**

And the homepage matters more than it has in years. LLMs cite you instead of linking you, so people type your name in and land on the front door.




PRESENTED BY



ILLUSTRATIVE

Traffic down does not have to mean signups **down**



Whatever traffic we get, it is on us to make the most of it.

How we run it at Pipedrive



THE SETUP

- One repo  shared with the whole web team
- Context and knowledge folders: ICP, brand, past decisions
- Skills for the repeating work
- Analytics, previous learnings and insights wired in
- Ideate and build tests end to end



WHAT IT CHANGED

- Idea to running test, fast
- We build it ourselves
- Not blocked on engineering
- More tests shipped, more learned

5x

TESTS LAUNCHED

~40%

WIN RATE

+9-10%

AVG UPLIFT FROM
WINS

What a context folder looks like

```
context repo

pipedrive-web-team-ai/
  CLAUDE.md
  .mcp.json
  context/
    ideal-customer-profile.md
    customer-problems.md
    2026-strategy.md
    brand-update-2026.md
    market-research.md
    paid-marketing-traffic.md
    ... 22 files
  knowledge/
    seo-knowledge.md
    ab-test-learnings.md
    content-style-guide.md
    design-system.md
    web-best-practices.md
    cognitive-psychology.md
    ... 19 files
  content/
    hypotheses/
    hmw/
    designs/
  .claude/
    skills/ 34
    buyer-journey.md
    competitive-landscape.md
    decisions-log.md
    gtm-initiatives-2026.md
    product-overview.md
    segments/
    how we do things
    the spine
    CMS, analytics, Jira, Figma
    what the team knows
    the work itself
    ab-tests/
    copy/
    web-qa/
    agents/ 7
```

Simplified. The real one has 22+ context files and 20+ knowledge files.



One test. End to end.

A real one. Not a demo I set up this morning.



PRESENTED BY




How it actually runs

- 01 Ideation – a customer problem becomes How Might We statements
- 02 Check – has my team already run this test before?
- 03 Hypothesis – written up properly before anything gets built
- 04 Design – exploring variations in HTML until I am happy with it
- 05 Build – straight into Optimizely via MCP
- 06 Launch – QA pass, checklists, metrics
- 07 Analyze – what we learned, and ten things to try next (yes, a Skill for that too)


BEFORE I SHOW YOU THE RESULT

Which one **won**?

A • ORIGINAL, THE CAPTERRA BANNER



B • VARIATION, THE STICKY FEATURE BAR



Hands up for A. Hands up for B.

VISIT TO TRIAL UPLIFT • AND A PRETTY NICE BUSINESS IMPACT €€€

Four tests. Four **wins**.

+11.05%

GETTING STARTED, EN

+7.17%

HOMEPAGE

+5.00%

GETTING STARTED, PT

+4.15%

GETTING STARTED, CRM

What we stopped **arguing** about

- 01 Which features to show – the context files already know
- 02 What it should look like – we explore until it feels right
- 03 When engineering can pick it up – it goes in via MCP
- 04 Whose opinion is right – the test answers that

WHY THIS MATTERS IN A BIG COMPANY

Most of the time does not go into building. It goes into aligning. We handle it ourselves now, off the right context, and we turn up with data instead of opinions. Life is easier.

Start with the context. Not the Skills.

Skills are the fun part. They are worth nothing if it does not know your brand, your rules, your customer.



PRESENTED BY



Every step is a Skill

- ✓ How Might We generator – a problem becomes testable ideas
- ✓ Learnings check – find out if your team already ran it
- ✓ Test hypothesis – write the test up before it gets built
- ✓ Schema markup audit – what your site actually publishes
- ✓ QA pass, page brief, test iteration, learnings log
- ✓ All of them free to copy

TO WRAP UP WHAT WE COVERED

What I actually use it for

pipedrive

AT PIPEDRIVE

- Conversion tests, ideation through analysis
- Publishing into the CMS
- QA before anything goes live

>veebermedia

ON MY OWN PROJECTS

- Building and rebuilding sites
- Schema, FAQs, author pages, llms.txt
- Tools and resources worth linking to

pipedrive

>veebermedia

GET THESE SLIDES AND SOME OF THE SKILLS AT

veebermedia.com/mos2026

Or find me in the hall, or just DM me on LinkedIn. I promise I will be friendly.

☐☐☐ And the rest of it lives at theaiweb.co

☑ SKILLS

☰ PLAYBOOKS

☞ MCPS

☰ STACK

💬 PROMPTS

linkedin.com/in/margusveeber



PRESENTED BY

