

The AI Press Release.

A field manual for announcements that machines actually read.

By David Berkowitz

Founder, High Caliber AI · Author, The Non-Obvious Guide to Using AI for Marketing

INTRODUCTION

The first reader is no longer human.

For thirty years, a press release has been a 400-word document written for a journalist with a deadline. That journalist has been replaced — or at least intermediated — by a model. Today, when news breaks, the first thing to read your announcement is an LLM running on behalf of an analyst, a buyer, an agent, or a search engine. The second is a human. The third may not exist at all.

This white paper is a field manual for the press release redesigned for that reality. It covers the principles, the structural changes, and a complete before/after example you can copy. It is short on purpose. The format change is not complicated. The implications are.

PART ONE

Eight principles of the AI-native press release.

01 — Structured over stylish

Style is overhead. Models reward structure: headings, lists, tables, JSON. Write the release as data first, prose second.

02 — Citeable, by paragraph

Every paragraph should be quotable on its own. If a sentence depends on the one above it for meaning, an answer engine will misquote you.

03 — Quotes are facts, not flavor

A quote is an on-record assertion by a named person. Treat it as structured data: speaker, title, organization, contact, verification status.

04 — Machine-first metadata

ISO-8601 timestamps with timezone. Stable canonical URLs. JSON-LD for the announcement type. No 'today.' No /v2.

05 — One canonical version

One URL per release. No PDF + HTML + wire copy with three different headlines. Models pick the version they trust most and ignore the rest.

06 — Atomic claims

Number every fact. C1, C2, C3. Numbers are easier for both humans and models to cite, and they make corrections trivial.

07 — Linkable evidence

Every claim links to a primary source — filing, datasheet, benchmark, repo. Models weight content backed by evidence higher in retrieval.

08 — Built to be ingested

Assume your release will be re-rendered by a model into a paragraph in someone else's article. Write the version you want to see come out the other side.

PART TWO

The mapping: every part of the old release, redrawn.

Below is a one-to-one map from the press release of 1995 to the press release of 2025. Same intent. New substrate.

TRADITIONAL	→	AI-NATIVE
FOR IMMEDIATE RELEASE	→	Stable canonical URL
"Today" / "Recently"	→	ISO-8601 dateline + timezone
Marketing headline	→	Atomic claim headline
Lede paragraph	→	TL;DR snippet block
Body paragraphs	→	Numbered atomic claims
CEO quote	→	Verified, structured quote
Boilerplate 'About'	→	JSON-LD payload
media@ inbox	→	Named contact + SLA

PART THREE

A complete before & after.

Same announcement. ACME Inc. ships a new AI inference appliance. First written the old way, then the AI-native way.

BEFORE · the 1995 way

FOR IMMEDIATE RELEASE

ACME Unveils Revolutionary New Platform That Reimagines the Future of Enterprise AI

SAN FRANCISCO — Today, ACME Inc., a leading provider of next-generation enterprise solutions, announced the launch of its most ambitious product yet — a paradigm-shifting platform that reimagines what's possible when cutting-edge AI meets best-in-class user experience.

"We are absolutely thrilled to embark on this exciting new chapter and partner with the world-class customers who share our bold vision," said Jane Doe, CEO and visionary founder of ACME.

About ACME: ACME is a leading provider of next-generation enterprise solutions empowering organizations worldwide to unlock the transformative potential of artificial intelligence at scale. Media contact: media@acme.com.

AFTER · the AI-native way

acme.com/press/reactor-2 · AI-native v1.0

ACME Inc. ships Reactor 2 at \$24,000, delivering 18,400 tokens/sec at 650W.

2025-05-06T09:00:00-07:00 · SAN FRANCISCO · en-US

TL;DR

ACME began shipping **Reactor 2** on May 6, 2025 at **\$24,000**, delivering **2.3x** Reactor 1's throughput at the same 650W envelope.

ATOMIC CLAIMS

- C1 · Ships 2025-05-06 at \$24,000 USD.
- C2 · 18,400 tok/s · Llama-3.1-70B fp8 batch=1.
- C3 · 650W sustained · MLPerf v4.1 protocol.

QUOTE · on-record · verified

"Reactor 2 hits 18,400 tok/s on Llama-3.1-70B at 650W. The same workload on Reactor 1 took 1,500W."

— Jane Doe, CEO · jane@acme.com · verified 2025-05-06

COMPARISON

Metric	Reactor 1	Reactor 2	Δ
Throughput (tok/s)	8,000	18,400	+2.30x
Power (W)	1,500	650	-56.7%
Price (USD)	\$28,000	\$24,000	-14.3%

STRUCTURED PAYLOAD · JSON-LD

```
{ "@type": "ProductLaunch",  
  "name": "ACME Reactor 2",  
  "launchDate": "2025-05-06",  
  "price": { "value": 24000, "currency": "USD" },  
  "metrics": { "tok_s": 18400, "watts": 650 } }
```

CONTACT

Mira Chen · Head of Comms · mira.chen@acme.com · +1 415 555 0142

Hours: 09:00–18:00 PT · Response SLA: 2 hours

PART FOUR

The 12-point checklist.

Use this before you publish anything.

- 01. Headline contains a subject, a verb, and at least one number.
- 02. Dateline is ISO-8601 with explicit timezone — no 'today.'
- 03. Canonical URL is permanent and version-free.
- 04. TL;DR is two sentences, structured to be the LLM snippet.
- 05. Every claim is numbered (C1, C2, C3...).
- 06. Every claim links to a primary source.
- 07. Quotes include speaker, title, org, contact, and verification status.
- 08. Comparison data is in a table with units and deltas.
- 09. JSON-LD payload is present, valid, and matches the prose.
- 10. Contact is a named person with role, hours, and SLA.
- 11. There is exactly one canonical version of this release.
- 12. An AI-summarized version of this release would still be accurate.

ABOUT THE AUTHOR



David Berkowitz is the founder of High Caliber AI and AI Marketers Guild, and he's the author of *The Non-Obvious Guide to Using AI for Marketing* (Ideapress, 2025). A longtime marketing strategist, David has led marketing and innovation for companies including Mediaocean, Storyhunter, Sysomos, MRY (Publicis), and 360i (Dentsu). He has contributed 600+ columns to outlets like Advertising Age, MediaPost, and VentureBeat, and spoken at 400+ events worldwide. He helps marketers harness AI to work smarter, stay creative, and strengthen customer connections. David lives in New York City.

Read more, get the templates, and join the mailing list at theaipressrelease.com.

Brought to you by High Caliber AI — highcaliberai.com.