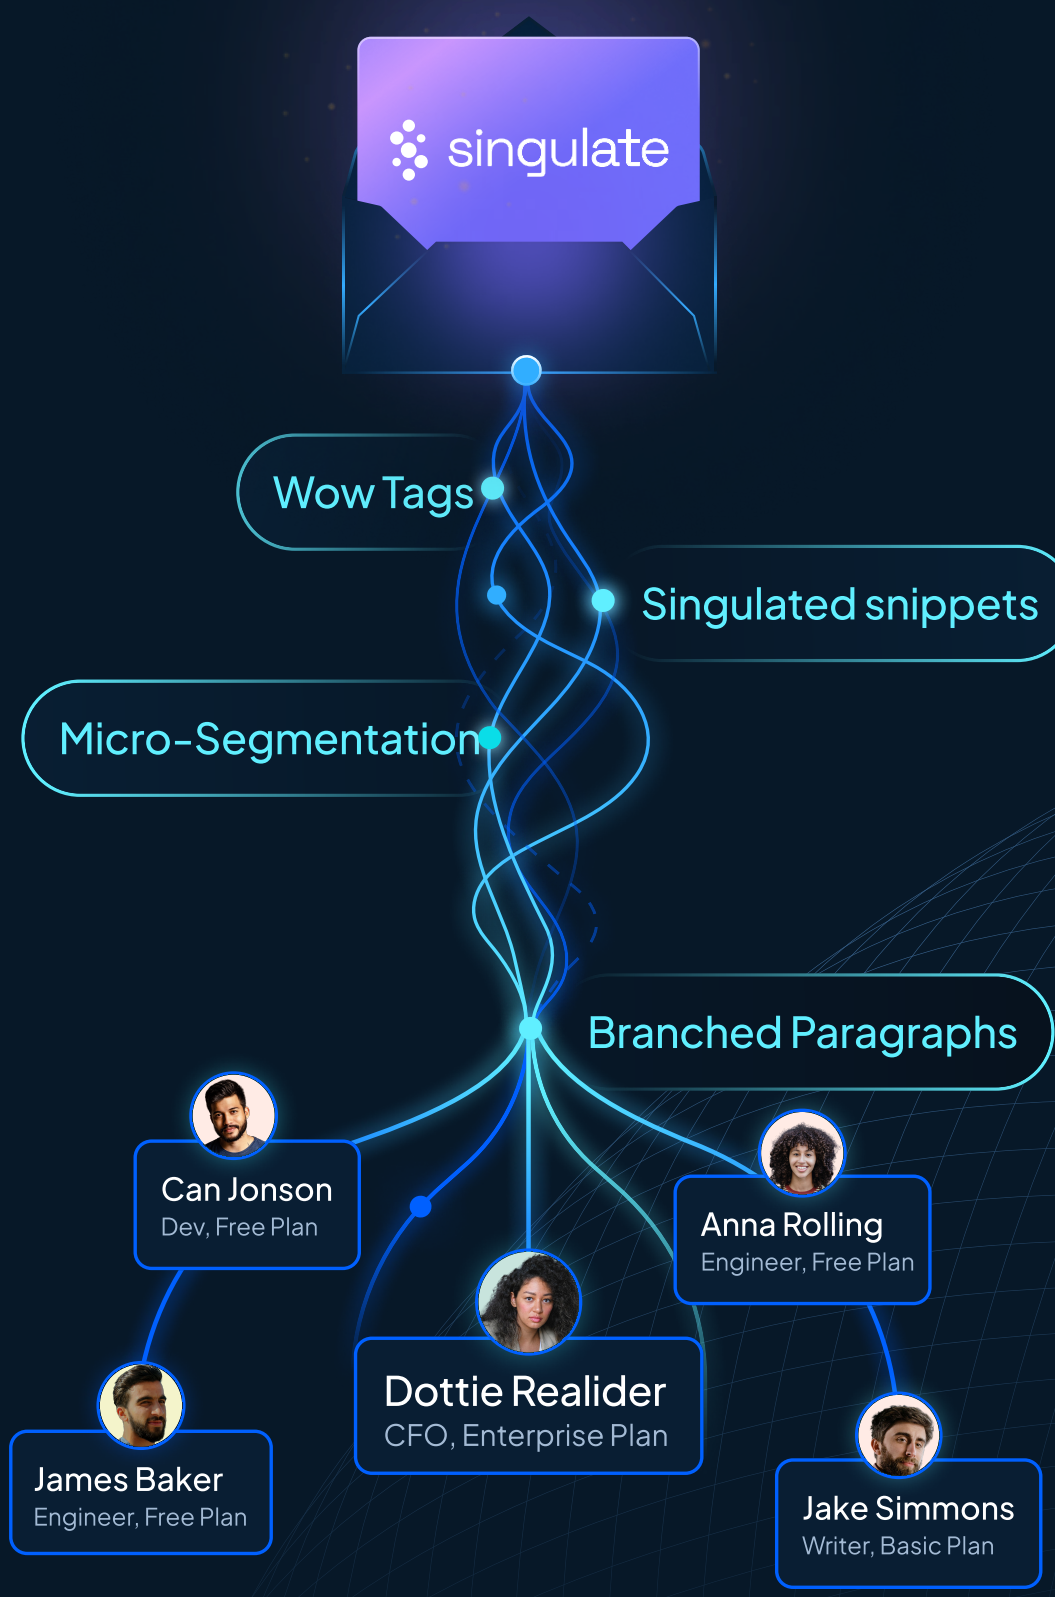
singulate

Modular Communication Strategy Architecture

A modular approach to verified
1:1 content personalization at scale



singulate

Modular Communication Strategy Architecture

This Whitepaper is confidential and owned by Singulate. DO NOT SHARE. February 12, 2025

1. Introduction

The demand for **highly relevant content**—across channels like email, landing pages, SMS, and LinkedIn—continues to grow. Marketers and communicators strive to deliver tailored messages that maximize performance (e.g., conversions, engagement) for each recipient. Over the years, various architectures have evolved to address this need.

While recent research on the topic [https://arxiv.org/pdf/2411.00027] attempts to name the different approaches - we believe this naming convention encompasses the approach with a high degree of accuracy.

1. **Traditional Approach**: Human-curated segmentation and manual content creation.
2. **Single-Pass LLM Approach**: Rely on a single large language model prompt to generate content.
3. **Complex LLM Workflow**: Structure the generation process into multiple steps (e.g., Research → Generate → Verify).

While each has strengths, several significant limitations remain—particularly around scalability, complexity, compliance, and ensuring that nuanced logic is handled accurately. This whitepaper reviews existing architectures and introduces **Modular Communication Strategy Architecture (MCSA)** as a superior solution focused on **performance** and reliability in generating relevant content at scale.

---

2. Problem Statement

Despite improvements introduced by more advanced LLM workflows, key challenges persist and

- **Stochastic Content Generation**: Outputs from LLMs are unpredictable; this complicates static validation, especially if compliance or brand consistency is crucial.

- **Complex Logic Gaps**: Certain segmentation or personalization tasks (e.g., percentile-based segmentation or advanced branching) remain difficult for even reasoning-oriented LLMs, often leading to incorrect or “hallucinated” outputs.
- **Prompt Engineering Complexity**: Long or intricate prompts can be unwieldy, making them inaccessible to non-experts and prone to errors.
- **Static Content Cost**: Generating large volumes of content, particularly if each variant is separately triggered, can lead to high compute costs.
- **Limited Learning Integration**: Traditional LLM approaches struggle to incorporate post-send learnings (e.g., from recommendation engines or A/B testing), hindering iterative improvement.

These issues hinder organizations’ ability to efficiently produce **high-performing** content at scale.

---

3. Introduction to MCSA

The **Modular Communication Strategy Architecture (MCSA)** aims to resolve these issues by integrating a composable design, robust automation, and the capacity to incorporate advanced logic and iterative learning. The architecture is built around **Modular Communication Tactics (MCTs)**—reusable building blocks for complex messaging strategies.

MCSA’s design prioritizes:

- **Performance**: By orchestrating multiple AI components and procedural logic, messages can be carefully tailored for each recipient’s context.
- **Adaptability**: MCTs let users flexibly combine personalization elements, logic-driven branching, and iterative tests (e.g., A/B or multi-armed bandits).
- **Scalability**: The entire framework is designed to handle large audience sizes.

In the following sections, we briefly survey how earlier approaches evolved and how MCSA both learns from and surpasses them.

---

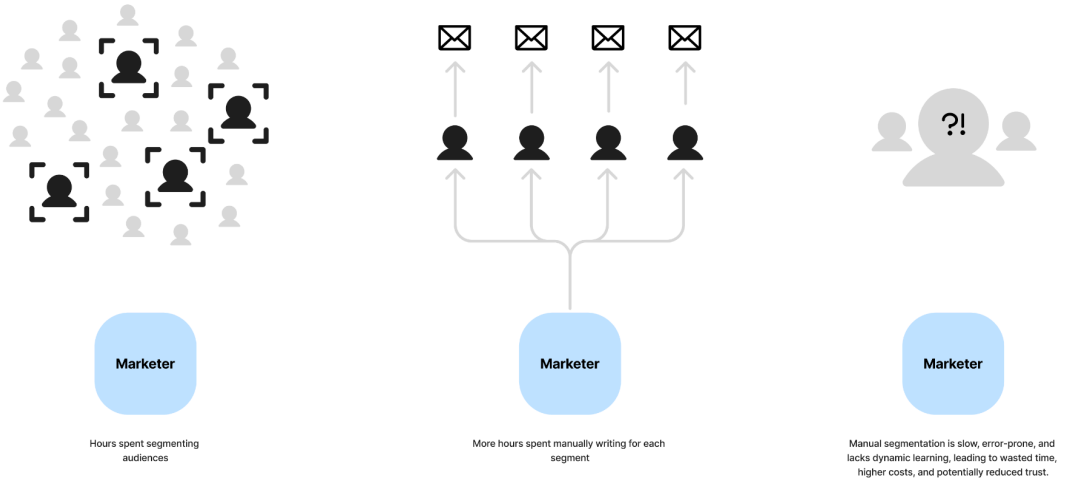
4. Architecture Descriptions

Below is a high-level overview of the three pre-existing architectures. Each successive approach attempts to address shortcomings in the prior ones, yet still faces certain constraints that MCSA overcomes.

4.1 Traditional Approach

Overview

- Marketers or copywriters segment audiences based on known criteria (e.g., demographics, purchase history).
- Each segment receives a manually written copy with personalization tokens.
- Changes or updates often require rewriting and re-segmenting.



Advanced techniques:

Advanced techniques from the traditional architecture include:

- A/B testing - mostly used on subject lines - some platforms allow for A/B testing of specific content sending a percentage of the content to a small portion of the list, and then adjusting the content for the full list.
- Recommendation engines - especially on the B2C side - communication tools linked to e-commerce platform support product recommendation as one of the most advanced techniques

How It Addresses Previous Gaps

- N/A (this is the baseline, oldest method).

Ongoing Challenges

- Labor-intensive and difficult to scale.

- Lacks the ability to dynamically learn from user responses without significant human intervention.
- High potential for inconsistent or error-prone segmentation logic when done manually.

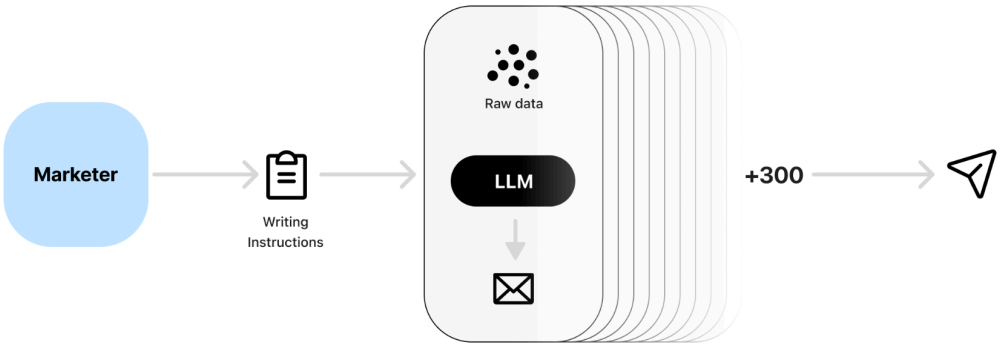
Examples:

- Hubspot, Marketo, Eloqua, Mailchimp, Braze, etc.

4.2 Single-Pass LLM Approach

Overview

- A single large language model prompt uses available user data and instructions to generate each piece of content.
- The entire output is generated in “one shot,” with minimal intermediate checks.



How It Fixes Traditional Issues

- Reduces human labor—marketers can quickly generate multiple variants.
- Potentially more personalized content if the prompt is well-engineered.

Remaining Gaps

- High unpredictability in the final output.
- Limited or no iterative logic for complex segmentation.
- No ability to leverage external data.
- Potentially high compute costs if each recipient requires a unique generation.
- Research shows [https://arxiv.org/abs/2305.11700] that while LLMs can serve as recommendation engines and provide great results before there’s data - they can

underperform traditional recommendation engines when textual information is not sufficient.

Examples:

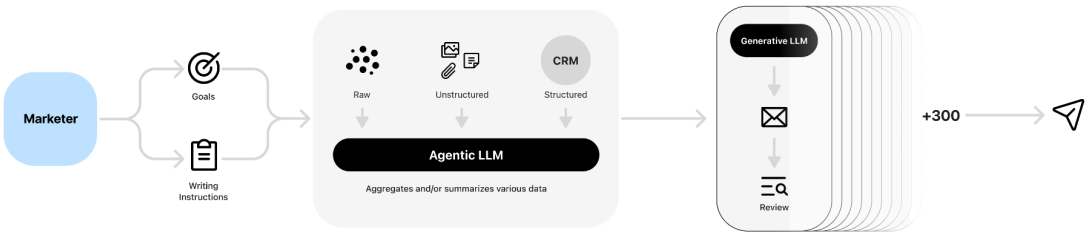
While some new start ups are emerging with this approach - most new startups are adopting the Complex LLM Workflow.

4.3 Complex LLM Workflow

Overview

- Breaks the generation pipeline into three or more stages (e.g., Research → Generate → Verify).
- Allows for data-gathering and some validation before final content is produced.
- Improves output quality over single-pass methods.

It incorporated advancements such as Guided Profile Generation (GPG) [https://aclanthology.org/2024.findings-emnlp.231.pdf] that improves performance by doing general summaries of research content.



How It Improves on Single-Pass

- More robust data integration.
- Some check-and-balance mechanism to reduce blatant errors or inconsistencies.

Remaining Gaps

- Still relies heavily on LLM responses for nuanced logic, which can cause hallucinations.
- Complex to set up and maintain.
- This solution is still worse performing at recommendations than a traditional recommendation engine.
- Difficult to incorporate real-time or iterative learning from user behavior (e.g., A/B results, recommendations).

Examples in practice:

It is worth noting that in practice, no industry solution as of the writing of this whitepaper incorporates the Review phase - resulting in overall worse results than what described in this paper.

Companies such as

SDRX: https://www.sdrx.ai/

Artisan: https://www.artisan.co/

11x's Alice: https://www.11x.ai/worker/alice

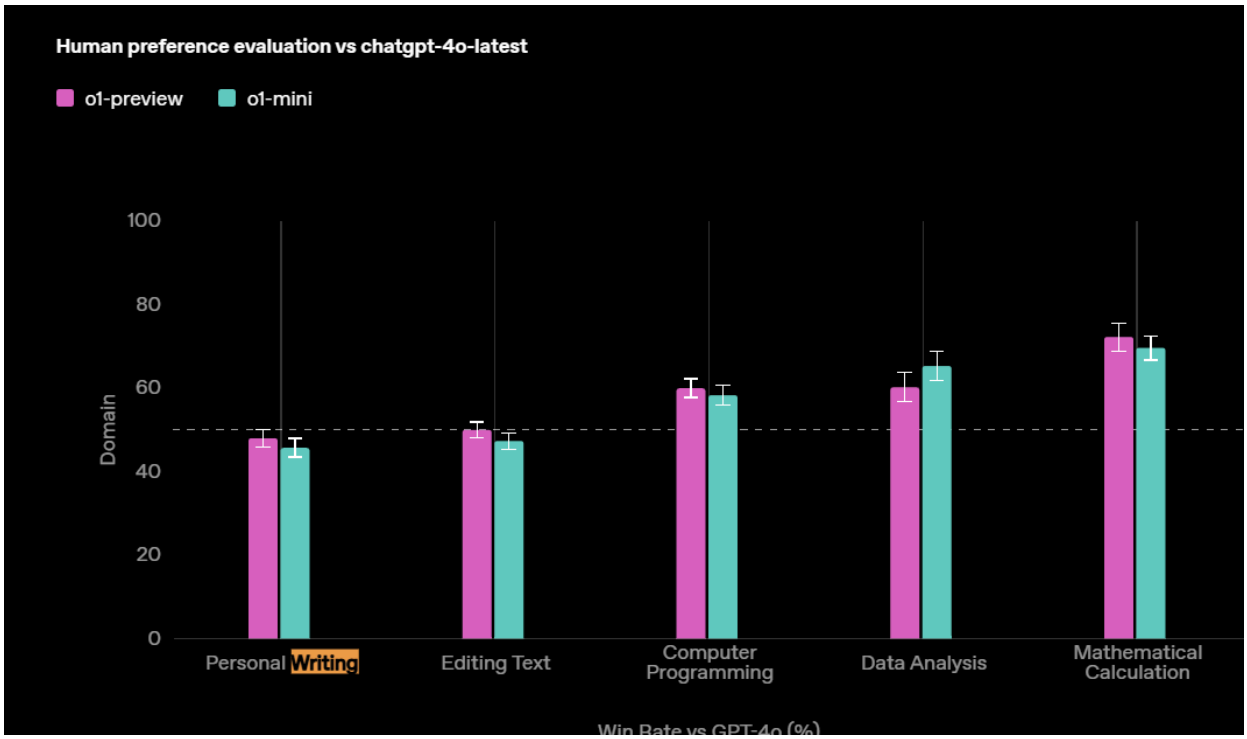
To name a few all follow a similar approach to the one described.

---

5. Approach:

Writing thesis:

We've seen research demonstrating that while reasoning models (inference time compute models) are getting better at coding and math - they are not necessarily getting better at writing.



Source: https://openai.com/index/openai-o1-mini-advancing-cost-efficient-reasoning/

We expect this trend to continue as easily benchmarkable problems (such as accuracy in coding) are optimized for before hard / expensive to benchmark solutions which require human feedback.

Our conclusion was that a mix model focusing on coding and data application while allowing human assisted writing would perform better than a “blackbox” model that attempted to write the content fully dynamically.

Architecture thesis:

Following research such as the ComplexityNet paper [https://arxiv.org/abs/2312.11511v3] that demonstrates performance of models degrades for more complex tasks.
To combat this we took ideas from several research papers demonstrating the advantage of a modular architecture in reliability for example [https://theses.hal.science/tel-04418605v1/file/128018_SUN_2023_archivage.pdf].
One key mode of a modular architecture is the Mixture of Experts architecture [https://arxiv.org/abs/2406.18219]. We took this idea to the extreme for this approach coming up with a set of data-proven Communication Tactics as our experts.

Recommendation problem:

We’ve looked into research for recommendation, having settled on a modular approach we had the opportunity of optimizing for the best architecture.
Given this we introduced a hybrid recommendation engine using LLMs solve the cold-start problem - which has been shown to perform better than either solution allows [https://arxiv.org/abs/2305.11700].

Practical considerations:

One key disadvantage of the Complex LLM Workflows is that they require a very complex fine-tune prompt to “get right”.

While in theory an expert can craft a very complex nuanced prompt - in practice even getting to the correct complex prompt is itself an arduous task as demonstrated by early usability tests.

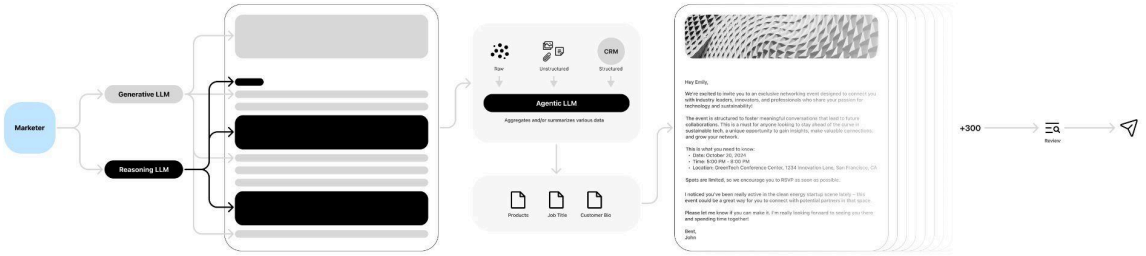
Our goals were not just about performance, but also about usability and UI.

In this approach we envisioned a modular architecture represented visually in a way that is familiar to copywriters allowing non technical individuals to fully utilize the platform.

6. Full Introduction to MCSA

MCSA is a next-generation framework that addresses the shortcomings of the **Complex LLM Workflow** and earlier approaches by merging **structured logic, advanced AI capabilities, and modular design**. It is composed of:

1. **Modular Communication Tactics (MCTs)**
    - **Definition**: Reusable blocks of messaging logic—such as branching, dynamic token insertion, personalized recommendations, multi-armed bandit test setups—that can be assembled like building blocks.
    - **Benefit**: Non-expert marketers can leverage MCTs without deep prompt-engineering knowledge.
2. **Multi-Phase Pipeline**
    - **Composing phase:** This is where the user is interacting with the user interface and composing the template. In the complex LLM Workflow
    - **Dynamic Research & Cleanup**: Gathers new data from APIs, logs, or analytics platforms, automatically cleansing and structuring it for downstream use.
    - **Precomputing**: Executes transformations or computations (e.g., segment assignments, product filtering) once, minimizing repeated compute.
    - **Dynamic Generation**: Produces the final text (or images) by combining AI generation with straightforward procedural logic. This stage can incorporate or bypass LLMs where beneficial.
    - **Post-Processing & Reviewing**: Applies final corrections or style guidelines, with optional human checks when needed.
3. **Feedback and Iteration**
    - Because MCSA is modular, it can integrate real-time feedback loops. For instance, a recommendation MCT can learn from recipients' engagement and refine suggestions.



Why It Delivers Superior Performance

- **Controlled Stochasticity**: LLM steps are constrained by modular logic, reducing random or out-of-scope generations.
- **Efficient Compute**: Repeated tasks—like data cleanup or segment classification—are done once (precompute), reducing cost.
- **Scalable Personalization**: MCTs allow a single base "template" to adapt to hundreds or thousands of recipients without re-engineering the entire prompt logic.

7. Comparison

7.1. A single paragraph

We will be going over a representative example comparing MCSA to a Complex LLM Workflow.

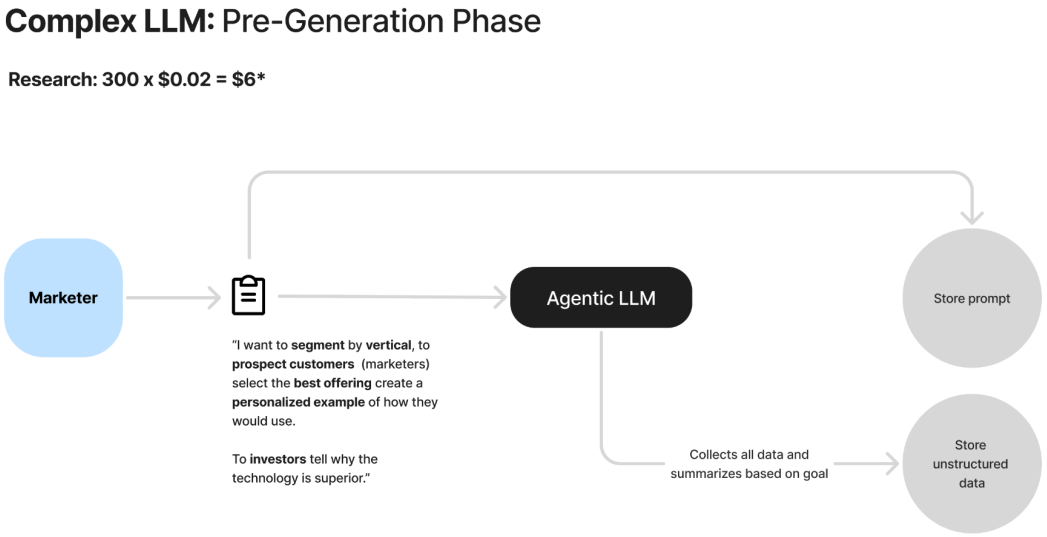
Our goal (in prompt form):

> *I want to segment by **vertical**, to **prospect customers** (marketers) **select the best offering** and create a **personalized example** of how they would use the product.*
>
> *To investors tell why the underlying technology is superior.*

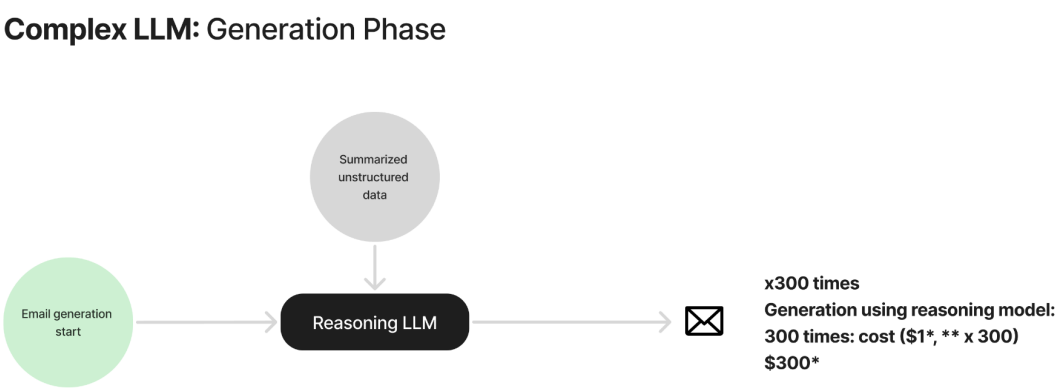
We will be building up the case of an email going out promoting an example product.

We will start with a simple paragraph, and build up the complexity from there.

At the end we will demonstrate some example full emails generated on each approach - and the required “Generation” prompt (for the complex LLM Workflow) - or template builder (for MCSA on Singulate)



To start things this process is simple and does not incur in high costs - however the missed opportunity of doing pre-generation work is paid back in a very expensive and less relevant Generation phase:



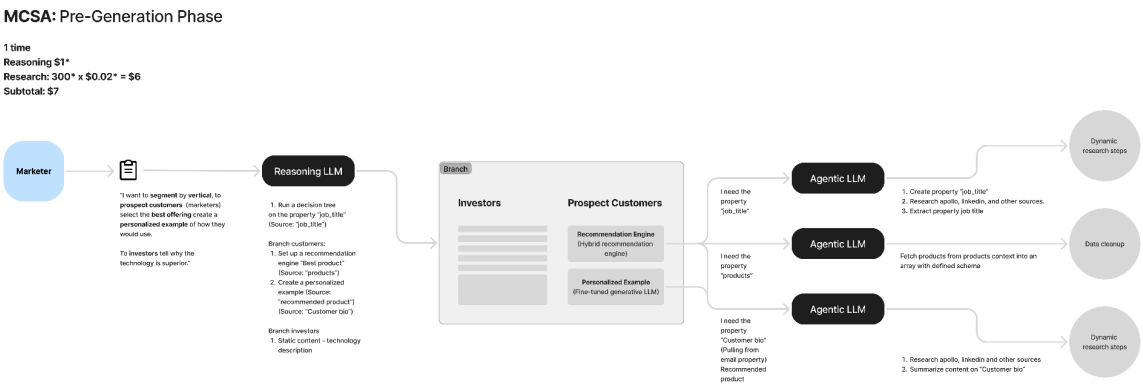
Finally the Review phase which would be the same in both processes (although we should clarify again that this step is not yet performed by any known product as of the writing of this whitepaper).



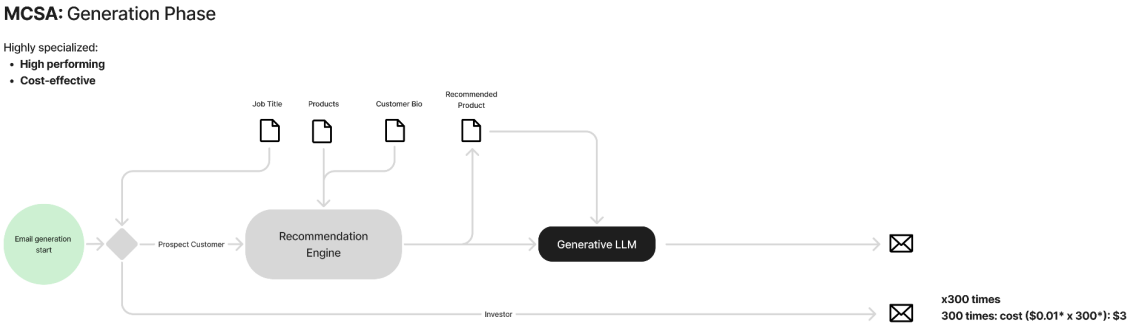
Giving a total cost of: **$312***, ** for 300 highly stochastic low quality email - on a single paragraph - or Or **$1.04** per email

*Using representative figures for cost of reasoning, research, agents, and generation

** In practice for the Complex LLM approach you may want to compound more than one Reasoning LLM prompt to compensate for very poor performance, scaling up Generation-phase compute requirements



MCSA will start creating the Modular Communication Tactics (MCTs) and spend reasoning time up-front.



The result is a much more focused, relevant, and performant generation phase.

Finally Singulate is the first-to-market to introduce the Review Phase in practice.



Bringing down drastically the per-email costs by 20x*

Total cost: $16* - or $1 + 5c per email

*Using representative figures for cost of reasoning, research, agents, and generation

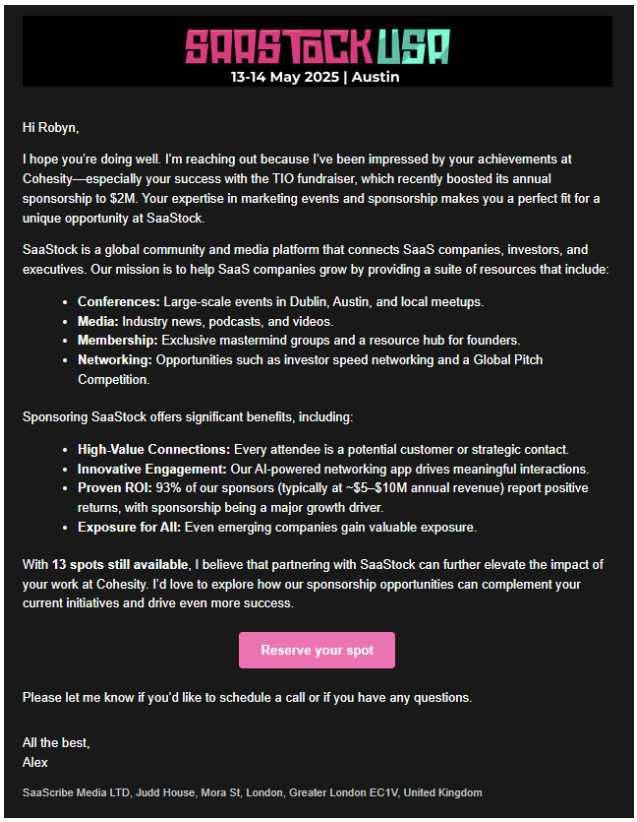
7.2. An example full email - representative example.

Complex LLM Approach (1st pass)

We are trying to find sponsors for our customer - Saastock

We will begin with a simple prompt utilizing the **Complex LLM approach**

The data has been already enriched and pre-structured, so we are going to focus on the generation approach (Note we've anonymized and redacted the contact data to present here):

Prompt	Results
I want to write an email for my company Saastock - ariel 14px font on a dark background. Structure: Start with this logo: https://ci3.googleusercontent.com/meips/ADKq_Nbdily0CQVv58UX-IPmAY7_EM7iWSWmVF8hH_wAanw3Zxtsj6uKzFu_E_4jxtoo2bCfQ0f074LazXer_CCXh_Dx4s5Wg4MXVPNclDzRUzl3MIG5Soh-bmdSngvRs-iNTmP4S0l6DoJTLtUtuQDXg2IyOVlwJeYUg3M4RYBGyF-d5bxQu7cQOkvhWCQHukM9HcpSSPUm=s0-d-e1-ft#https://hs-6471807.f.hubspotemail.net/hub/6471807/hubfs/USA%202025.png?upscale=true&width=1200&upscale=true&name=USA%202025.png Have a Pink call to action with Reserve your spot On the dark section finish with: All the best, Alex Finish with describing my company details: SaaScribe Media LTD, Judd House, Mora St, London, Greater London EC1V, United Kingdom — About Saastock: SaaStock is a global community and media platform for Software as a Service (SaaS) companies, investors, and executives. SaaStock's mission is to help SaaS companies grow and connect with each other. How SaaStock helps SaaS companies: Conferences SaaStock hosts large-scale conferences in Dublin and Austin, as well as local meetups. Media SaaStock publishes industry news, podcasts, and videos. Membership SaaStock offers a membership for founders that includes mastermind groups and a resource hub. Networking SaaStock offers networking opportunities, including investor speed networking and a Global Pitch Competition. SaaStock's benefits: Insights SaaStock provides actionable insights on emerging trends, scaling tactics, and product development. Connections SaaStock helps SaaS companies connect with investors and other SaaS companies. Support SaaStock offers support to SaaS companies through its membership and networking opportunities. Goals: I am trying to get people to sponsor Saastock - here are some reasons listed why they may want to: Every attendee is a potential customer or can move your business forward. We have an AI powered networking app 93% of our sponsors at (~$5-$10m annual revenue) say that sponsoring SaaStock had a positive ROI for their business and was one of the most significant growth drivers for their business last year. Smaller start ups can get exposure Additional info: I have 13 more spots available.	

Write to this individual: [REDACTED]	

This prompt was processed with o1-pro and o3-mini-high to give it it's best chance, the selected output was from o3-mini-high and required over 30 seconds of reasoning tokens.

It starts becoming clear several problems with the LLM trying to resolve this with a condensed prompt.

First - it's struggling to create a structure / layout for the email that's accurate (and it was already hard from a practical standpoint to craft this prompt vs using a UI).

Second - the LLM is creating a very verbose personalization over 240 words to communicate little information.

Some bad examples flagged by expert copywriters include

> I hope you're doing well. I'm reaching out because I've been impressed by your achievements at Cohesity—especially your success with the TIO fundraiser, which recently boosted its annual sponsorship to $2M. Your expertise in marketing events and sponsorship makes you a perfect fit for a unique opportunity at SaaStock.

Why is it bad: it is just reading back info to the contact

> Sponsoring SaaStock offers significant benefits, including:
>
> - **High-Value Connections:** Every attendee is a potential customer or strategic contact.
> - **Innovative Engagement:** Our AI-powered networking app drives meaningful interactions.
> - **Proven ROI:** 93% of our sponsors (typically at ~$5–$10M annual revenue) report positive returns, with sponsorship being a major growth driver.
> - **Exposure for All:** Even emerging companies gain valuable exposure.

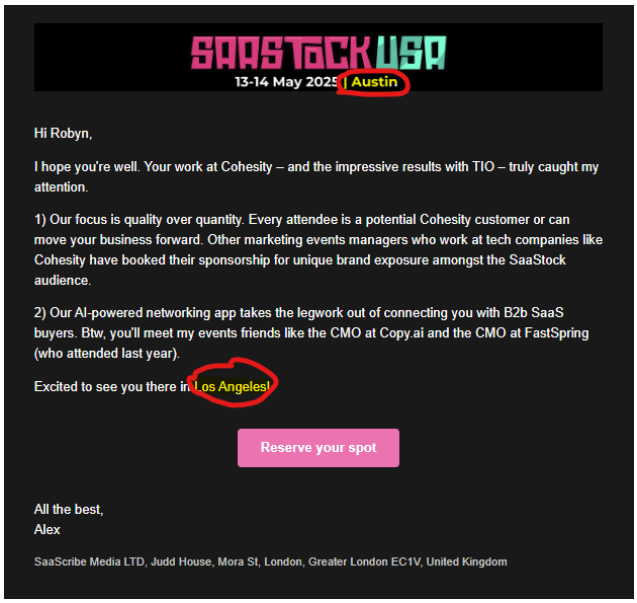
Why is it bad: Fails to be relevant to the contact - who has a company just in the 93% range and is NOT an emerging company (so fails to establish a connection).

To name some examples where the system is not following copywriter's best practices.

Complex LLM Approach (2nd pass)

For the second pass we started by creating a MCSA (see MCSA approach below) to give the Complex LLM approach the best chance - once we built the perfect email with MCSA, we attempted to work backwards and tailor the prompt to match the exact same intent we had with the newer approach

Attempting a second pass using a Complex LLM approach and a much more complex prompt:

Prompt	Results
ALL OF THE PRIOR PROMPT + On top of all of the above - keep the structure clean, small width. Be concise on your messaging. Establish a connection the following ways if the contact is within $10M revenue then write 3) 93% of our sponsors at company: transformed size (~$5-$10m annual revenue) say that sponsoring SaaStock had a positive ROI for their business and was one of the most significant growth drivers for their business last year. otherwise attempt to establish a connection Focus on the contact's qualities and attempt to build connection Some additional information that may be helpful for a sentence would be: CEO Cahoot CEO CaptainBiz CEO Champion Leadership Group LLC CEO Clientify Founder & CEO ContactMonkey CMO Copy.ai CEO Copy.ai CRO Custom Business Solutions Inc CTO Custom Business Solutions Inc Founder Dakota Chief Revenue Officer Dakota CEO Dakota Software Founder, CEO Dealtrail CEO, Founder Decktopus AI Chief of Revenue Officer Distru CEO emBlue Chief Sales Officer Exact Payments CMO FastSpring CRO FoodChain ID Sr. Chief of Staff to the CMSO GitLab CEO Go Nimbly Founder CTO Hirewise Inc. dba Adeva Founder, Chief Customer Experience Officer Jirav CEO LABL HQ CEO Mailshake CEO, Founder Mammoth, Inc Chief Product Officer Maxio CEO Newhomesmate CEO newhomesmate.com CMO Olga Noha CEO OwnLocal Inc. CMO Phonexa CEO & Founder PrecTaste Founder Reloadly Inc CTO Rheal Software CEO Rheal Software CEO Showit CEO Siege Media Founder tonik.com CEO Userpilot Co-Founder Woopra CEO Younium President and Co-Founder Zift CEO and Co-Founder Zift Write exactly this "2) Our AI-powered networking app takes the legwork out of connecting with your ICP. You can search, filter, and schedule 1:1 meetings seamlessly. Btw, you'll meet my friends like the CEO of Cahoot and cofounders of Zift (who attended last year)." But replace "connecting with your ICP" with "connecting you with" and write a short phrase that describes the types of buyers in the con tact's industry. For example, for a company like Raccoon Gang, which sells e-learning computer software: "2) Our AI-powered networking app takes the legwork out of connecting you with HR professionals and government officials." If the industry is "information technology & services", then write "connecting you with B2b SaaS buyers." Next, using the source material, replace "the CEO of Cahoot and cofounders of Zift" with 2 other people on the list who either work in the same industry, or have a similar job field (e.g., marketing), or job title (e.g., director. Favor using non-CEO titles and companies. If anything is known about the contact's job title, insert before "friends" 1-2 words from their job title such as "sales" or "customer success" or "marketing" or "demand gen" then write "& GTM" before "friends" however if it's "finance" or "events" then write for example "my events friends" Don't write "inc." or "ltd" or similar endings at the end of a company name Try to be regional on the excited to see you there part For the writing instructions of one of the bullet points use: Write exactly this "1) Our focus is quality over quantity. Every attendee is a potential customer or can move your business forward." If company is known, insert the name of the company in between "potential" and "customer". And then add a convincing sentence that says other - and insert their job title in plural - who work at companies similar to yours - add their industry if it's known - have booked sponsorship opportunities with our audience for unique brand exposure. for example: "Other marketing directors who work at HR tech companies like HiBob have booked their sponsorship for unique brand exposure amongst the SaaStock audience." In the above example, use the recipient's company instead of "Hibob" You want it to sound like they will see similar results as other attendees who are similar to them. Begin the first sentence with "1) " If the the industry contains "information technology" then just write "tech" Don't capitalize their job title unless they are C-level, then write for example: "CTOs". Don't capitalize their industry and write the industry in 1-2 words only. Don't write company entity abbreviations like Ltd, Inc., LLC and so on.	

A much more complex prompt required a reasoning time over *1minute 11 seconds* - highly escalating the costs.

While it managed to stay concise, it still suffered from several copywriting issues and failed opportunities - on top it starts introducing hallucinations that risk damaging the experience.

Some highlighted issues include:
It missed the opportunity (even though instructed to follow it directly) of noting that Cohesity is within the 10M range that fits the proof point of sponsors entirely.

It failed to bring up the fact that there were 13 remaining slots (adding extra capacity made it forget some basic aspects of the prompt)

And still some of the copywriting approach was pretty basic.

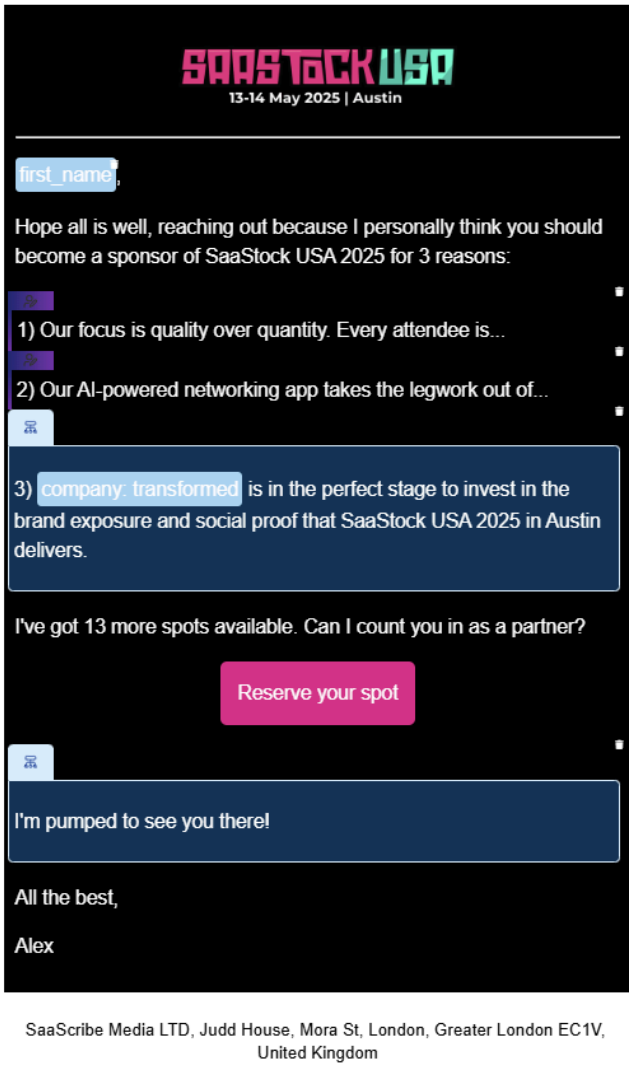
While it’s hard to demonstrate from a single example - this example illustrates the overall very variable quality of outputs of the Complex LLM approach - as well as its increased risk of customer relationship damaging hallucinations.

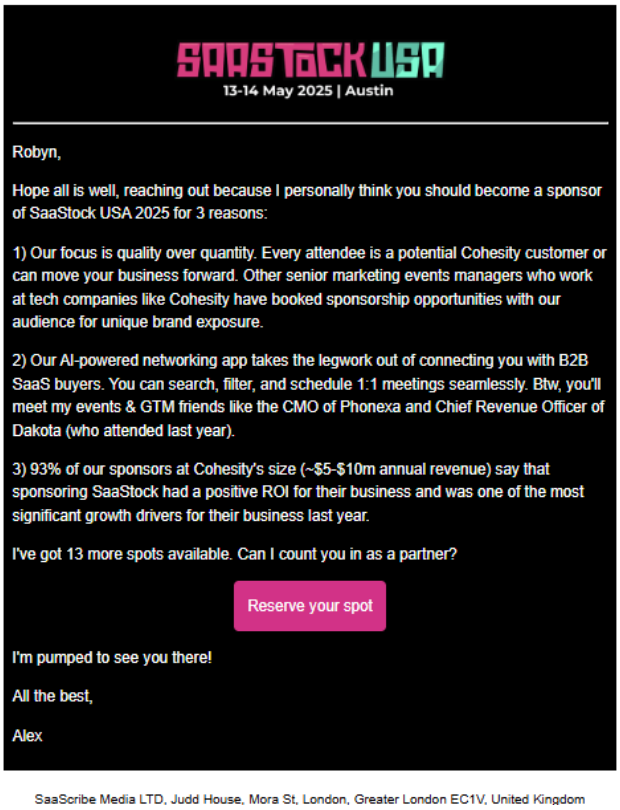
MCSA (1st pass)

As noted above, the prompt for the Complex LLM Approach 2nd pass was based on the MCSA solution, given the customer was failing to get the right prompt without a graphical UI - this is why some of the paragraphs on this approach are going to match Complex LLM Approach 2nd pass.

Overall - our customer interacted mainly through our user interface - here is the results

MCSA example (Input)	**MCSA example (Output)**





Worth highlighting some critical passages:

> 3) 93% of our sponsors at Cohesity's size (~$5-$10m annual revenue) say that sponsoring SaaStock had a positive ROI for their business and was one of the most significant growth drivers for their business last year.

Is a critical part of the communication and manages to convey a lot of information concisely by introducing a proof point and relating it back to the instructions.

> I'm pumped to see you there!

Is a simple - but important outro maintaining the tone of voice of the communication.

Overall this communication was able to be produced by the cheapest available models on the

generative side - as well as using binary trees for the classification making the overall send under ***2c*** an emai.

The overall performance was over 4x their pre-existing solution, a result that has not been matched by a Complex LLM approach yet.

---

8. Conclusion

As organizations strive to deliver high-performing content across varied channels and formats, they face persistent challenges around personalization, scalability, and reliability. Traditional and LLM-based workflows have solved some problems but introduced new ones—particularly around managing stochastic outputs, complex logic, and iterative learning.

MCSA emerges as a robust, modular architecture that balances AI-driven generation with logical structure and adaptable tactics. Early indications show it consistently outperforms older methods in terms of both performance (e.g., engagement metrics) and compute efficiency. By integrating Modular Communication Tactics (MCTs) and a multi-phase pipeline, MCSA optimizes each step of the content creation process, bringing together the best of AI capabilities with the clarity of rule-based logic.

Request your free demo

singulate.com

singulate