



The Journey from the Analog to the Agentic Chief of Staff


Hans Husmann


DOI: <https://doi.org/10.66241/jaaj6>



Hans Husmann

Scale AI





The Journey from the Analog to the Agentic Chief of Staff

by Hans Husmann

May 2, 2025

ISSN: 2995-5971

Abstract

This article ties the origins of the Chief of Staff role, its essence, and strategic importance to the possibilities of augmenting the work with GenAI solutions. Beyond a commentary on the technical implications for the role, it also gives a perspective on how to implement these systems in daily workflows. It further attempts to predict where the evolution of the underlying technology could take the role.

The median tenure of an S&P 500 CEO today stands at less than five years. A decade ago, this number was ~20% higher^[1]. Like their predecessors, today's corporate leaders contend with numerous internal and external problems on a daily basis. Yet in the present they do so within faster technological advancement, shorter news cycles, and a more interconnected world that regularly turns global crises into local ones. These mounting pressures and globally distributed teams have outpaced the capacity of a single executive to lead effectively. As one solution of managerial leverage today most enterprises employ a Chief of Staff. Jamie Dimon famously runs America's largest bank, JP Morgan, with a physical follow-up sheet. That story neglects that he has had his chief of staff, Judy Miller, and an army of executive assistants for over 20 years to turn "to-dos" into actions.

In theory, the chief of staff is tasked with giving leverage to a leader. In practice, that means helping to define the strategy and ensuring consistent execution of the

CEO's strategic priorities. Hereby the chief of staff inherits the obstacles of their principle: a sluggish flow of information, different and often inconsistent data, subjective biases, human error, and organizational inertia. In short, the kind of problems that Generative Artificial Intelligence (GenAI) technology is increasingly capable of addressing.

History of the role

The role of the chief of staff originated in the military, where they served as aides to the commander-in-chief (cf. Alexander Hamilton's effective chief of staff role to George Washington in the American Revolutionary War). Following the Prussian military historian and original strategist Carl von Clausewitz, their objective was to be the bridge in the leadership command of the Auftragstaktik^[2]. In this leadership philosophy, the military leader defines the objectives of the order—the commander's intent—which are in turn to be interpreted by the subordinate with the freedom to execute in a decentralized mode. The commander's aide (the de facto chief of staff) helps to both inform the initial strategy as well as monitor (and unblock) its execution.

After its military origins, the role became equally formalized and famous in the political definition of the White House Chief of Staff^[3]. This non-political appointee was formally created by President Eisenhower when he created the Executive Office in 1953. The formalization of the role at the centre of the most powerful post-war democracy is not random. First, new communication methods increased the speed of decision-making from days to hours. Secondly, the coordination effort in democratic institutions—that rely on consensus building and multi-threaded expert input—is far higher than in aristocratic or authoritarian models of government. The general public took until the television version in Leo McGarry of *The West Wing* or Doug Stamper from *House of Cards* of the "fixer role".

The role's rise in Silicon Valley coincided with the dot com boom and its explosion of high-growth startups centered around internet-based business models funded by growth oriented venture capitalists. Here the role was often introduced by the investors, who came to see it as leverage for the founder and a tool to bring execution discipline to their young and scaling investment. The position offered

proximity to leadership and exposure to strategic decisions. Over time, it evolved into a training ground for future founders and general managers.

The Chief of Staff role and its relevance today

Today, the role varies across three different layers of responsibility: from the elevated executive assistant, a young professional with a consulting/MBA background in a 1-2 year stint, to seasoned executives who serve as long-term right hands to Fortune 500 CEOs. My personal organizational definition as a Chief of Staff at a Silicon Valley artificial intelligence company breaks it down into three sequential areas of responsibility: shaping, executing, and telling the company's strategy (or story).

- **Shaping** - The definition of the strategy for the company, as set by the leader and executive team with input from across the organization. Here, the CoS helps to draft the strategy, gathers expert feedback, conducts market analysis, and helps to align the executive team.
- **Executing** - This part is often described as “running the trains on time”. It is keeping the ratio between what an organization said it *would do* vs. *what is being done* close to one. Practically, that means setting up effective systems of tracking and decision-making that keep a high velocity of (useful) information between the different teams and executives.
- **Telling** - Strategy here is purposely defined as a story, as human cooperation at its core relies on shared stories. Communicating strategy/executions and the underlying motivation is ultimately a method of giving leverage to achieving the objectives of the enterprise.

To fulfill the breadth of responsibilities above with the necessary depth, the chief of staff needs to build a system that makes their job manageable. Here most CoS operate within the setting of management systems that have evolved with the rise of the knowledge economy in the past 50 years: From Peter Drucker's foundations of modern management, to David Allen's Getting-Things-Done task management method, to Google's OKRs^[4] to - the author's underlying philosophy - management operating system as defined and coached by Silicon Valley's leading coach Matt Mochary^[5].

Even with the perfect productivity system and extreme discipline, it is nearly impossible to maintain this system without extreme long hours, stoic discipline and enormous context switching. Even with a perfect system, shortcomings on either the operational or strategic fronts of the role are inevitable. This changed fundamentally with the introduction of artificial intelligence, and specifically the availability of large language models and agents.

Building the Agentic Chief of Staff

The administrative “keeping the trains running on time” remit of every CoS is the textbook definition of recurring, high leverage but often highly manual tasks. This ranges from updating the weekly business review document with KPIs, to generating meeting follow-ups, and reviewing budget requests. All of these tasks are united by the need for written communication, a level of standardization in the documentation, clear owners, and regularity that drives accountability - necessities to ultimately progress on the organization’s objectives.

In essence, they are a perfect set of tasks to be augmented or fully automated via large language models. Before getting into the details on how this is best done, consider why this is relevant beyond time savings. Most, if not all, “human” chief of staffs suffer from the undoing of a knowledge worker: inconsistencies in execution, the tendency for factual errors missed by oversight in reviews, and becoming a bottleneck as the information does not flow through the organization fast enough. To resolve this, the GenAI augmented CoS solution needs to be built in three steps: starting with a clear definition of the tasks to be automated, through the technical implementation and testing, and finally the human-in-the-loop feedback cycle.

The process of building an agentic support system starts with process mapping of recurring tasks. Tasks fall typically into the areas of data gathering, analysis, interpreting, summarization, and communication. Across these, it is crucial to define the data sources, identify the directly responsible individuals, and clarify the format and cadence of communication. Here the chief of staff needs to become the product manager of their “management product” and synthesize the information in a product requirements document (PRD). The PRD serves as the ground truth on what the solution needs to solve as input to the LLM as much as

an objective benchmark against a quick fix automation. The PRD is then translated into a prompt, which initiates an iterative testing process.

For testing, it is best to start by selecting a small subset of the process, e.g., getting one data table with sales for one weekly report, before attempting to create the whole report in one go. The system relies on what I call protocols. A protocol is a simple checklist that captures the instruction for the model or agent as context in the prompt. This includes a step-by-step instruction on what steps to take, and then a list of checks that need to be done on the data for validation (e.g., formatting guidelines, typical errors to flag). In the testing phases, the input data purposely needs to include false flags to figure out if they are flagged through your protocol guidance. The model output should be refined iteratively by applying incremental prompt adjustments until the desired result is achieved. If multiple tweaks are made via additional prompts it is advised to request the model to add additional instructions to the original prompt. Based on my deployed solutions, the accuracy rate here is often at 80-90%.

AI driven software/feature development shifts the ownership of the tool development from the scarce resource of a software engineer and long development cycles to the manager. Typical avenues at most companies are enterprise licenses of OpenAI, Anthropic or the Gemini integration via Google Cloud. Given the often highly sensitive data being processed, like sales figures or strategic decisions, all of the solutions need the security team's review. The appeal of this form of 'vibe coding' lies in the ability of non-technical users to deploy solutions within hours. While this is remarkable to get started, it is not uncommon to see inspiring demos, but unusable solutions. They fail because underlying data sources change over time, or the model hallucinates to show completeness over accuracy - not dissimilar to human errors caused by wrong incentives. The negative scenario could be a version of Goethe's poem "The Sorcerer's Apprentice". In the poem a young magician requests magic - he cannot fully control yet - to take over his mundane cleaning tasks. Albeit with good intentions, his attempt to avoid the hard work, culminates in the iconic scene of multiplying broomsticks causing an uncontrollable flood. The enterprise version of that could be daily distributed pipeline updates taking a summary version from call recordings that do not differentiate business critical from personal chatter, perhaps not too unfamiliar to some readers.

A human-in-the-loop - in the form of the chief of staff - is what defines really effective drivers of this system: Effectiveness in execution depends both on the consistency of follow-ups and the quality of their curation. The curation is fine-tuning communication through feedback on the output of the LLM, and evolving instructions as the systems get deployed against the organizational realities. Communication here itself can be seen as a feedback loop, as I have witnessed that consistent, detailed communication generated through an LLM may lead to reduced engagement, as such messages are perceived as automated.

From agentic Chief of Staff to an agentic management system of “trial and error”

The positive vision however of this evolution is a mostly autonomous management system: A chief of staff becomes the “conductor” of all decision making and information flows for an enterprise leadership at a speed and quality not seen before. On a philosophical level, I imagine the role’s future in the tradition of Karl Popper’s idea of human progress through “trial and error”^[6]. Popper stated that human knowledge advances only through the refutation of initial hypotheses. I foresee technology-enabled leadership in the tradition of rational progress through error corrections. In a world where AI shifts its input-to-output mechanic to agentic, iterative processing— and where the human leader provides intent and feedback—the agentic chief of staff can become (an even more central role) in an organizational system of rapid learning and improvement.

About the Author



Hans Husmann

Chief of Staff at Scale AI

Hans Husmann is a strategy and enterprise business transformation professional whose work focuses on aligning organizational strategy, operating models, and technology-enabled change across industries with a global footprint. His professional interests center on the human-in-the-loop application of advanced technologies, including the governance and adoption of AI-enabled decision support in complex enterprise settings. Drawing on cross-sector experience, he has contributed to large-scale transformation initiatives that integrate process redesign, stakeholder alignment, and measurable value realization. Husmann's perspective emphasizes sociotechnical design—treating data, workflows, and human judgment as interdependent components of effective digital transformation. He is particularly interested in responsible deployment practices that improve transparency, accountability, and user trust in automated and semi-automated systems. Through advisory and leadership roles, he continues to advance pragmatic frameworks for translating emerging technology capabilities into durable organizational outcomes.

[LinkedIn](#)

Copyright Notice

All articles are published in the Journal of Business and Artificial Intelligence under the Creative Commons 'CC BY' ("Gold Open Access") license, where authors retain the copyright of their articles. The author grants JBAI a license to publish the article under a Creative Commons 'CC BY' license, which allows the work to be freely accessed, shared, and used under certain conditions.

About the Journal

The Journal of Business and Artificial Intelligence (ISSN: 2995-5971) is the leading publication at the nexus of artificial intelligence (AI) and business practices. Our primary goal is to serve as a premier forum for the dissemination of practical, case-study-based insights into how AI can be effectively applied to various business problems. The journal focuses on a wide array of topics, including product development, market research, discovery, sales & marketing, compliance, and manufacturing & supply chain. By providing in-depth analyses and showcasing innovative applications of AI, we seek to guide businesses in harnessing AI's potential to optimize their operations and strategies.