


Re-Engineering the Content Supply Chain

From Linear Content Manufacturing to
System-Centric Content Operations



Executive Summary

The job of a CMO has changed faster than the production system underneath it.

More markets, more SKUs, more channels, more variants, more compliance regimes, delivered by a flat or shrinking team, judged on growth rather than output. Brand integrity, speed to market, marketing ROI, and content effectiveness all sit on the CMO scorecard. All four break under the production model that supplies them.

That model still operates as a linear manufacturing process driven by humans. Scope, brief, ideation, hero content, review, localization, delivery. It was rational in a pre-digital world. It is now structurally incompatible with the speed, scale, and regulatory complexity of modern marketing, and incapable of producing the content effectiveness CMOs are actually measured on.

What is needed is not a faster version of the same production system. It is a re-engineered one: the brand owns the strategy of the system, and the system produces content that learns from market signals and improves over time. Efficiency is a by-product. Effectiveness is the destination.

This is a CMO-led transformation, not a content operations problem to delegate downward.

Contributors:

Micha Savelsbergh

CPO, INDG/GRIP

David Parker


Former VP Digital & Consumer Experience Design, PMI

AI layered onto a broken operating model only accelerates chaos


AI is often positioned as a solution to this problem. While this is theoretically true, so much of it depends on implementation. A rush to “AI everything,” in practice, can accelerate dysfunction.

AI without proper workflows can create increased busy work and actually decrease productivity. Faster production without systemic control increases risk, inconsistency, and governance overhead.

OLD MODEL




NEW MODEL



“You never change things by fighting the existing reality. To change something, build a new model that makes the existing model obsolete.”

Richard Buckminster Fuller Jr.,

American architect, systems theorist, writer, designer, inventor, philosopher, and futurist.



Fuller was most famous for his lattice shell structures – geodesic domes, which have been used as parts of military radar stations, civic buildings, environmental protest camps, and exhibition attractions.

This whitepaper proposes a fundamentally different approach: a **system-centric content supply chain** in which content is executed by design, not assembled by hand. In this model, design systems, data feeds, configurators, and digital product twins replace briefs, reviews, and offshore production lines. Compliance is embedded upstream, not enforced downstream.

The system is designed from the market to the center.

TABLE OF CONTENTS

01 The Proliferation of Content. ↓	06 What Disappears in This Model. →	07 Commercial Impact. ↓
02 Structural Failures of the Linear Content Supply Chain. ↓	05 What Organizations Must Change. ↑	08 Where This Goes. —
03 A Transformation is Necessary. →	04 The System-Centric Content Supply Chain Architecture. ↑	

01




The Proliferation of Content

Advertisers and marketers are under pressure to produce more content than ever.



A 2025 Adobe survey found that 96% of marketers saw content demands expand at least 2X in the past two years, with 71% of respondents believing that demand for content will grow 5X by 2027.

There are several pivotal factors for this exponential growth.



Globalization

In addition to translation requirements, brands also need native experts to keep them mindful of any local customs or traditions to either uphold or avoid offending.

Personalization expectations have risen sharply

Consumers expect tailored interactions, which pushes brands to produce many more versions of messages, offers, and creative. A McKinsey study found 71% of respondents expected personalized interactions from companies, and 76% expressed frustration when those companies failed to deliver. That same Adobe survey found that 61% of marketers cite consumer expectations for personalized experiences as the top factor driving content demand.

A McKinsey study found
71% of respondents expected
personalized interactions from
companies, and **76%** expressed
frustration when those
companies failed to deliver.

There are more surfaces and formats than ever before

On Google alone, there are opportunities to feature content on YouTube, Display, Search, Discover, Gmail, and Maps. Google explicitly tells advertisers to upload as many text, image, and video assets as possible because more assets let the system create more ad formats and appear on more inventory. In addition, TikTok, X, Threads, LinkedIn, Instagram, CTV, digital out of home, and more are all options for many brands. Google has specifically told brands that different combinations of content and messaging resonate differently with different users and queries, and that “asset flexibility” helps its systems assemble the best mix for each context. That means advertisers are pushed to create multiple headlines, descriptions, images, and videos so the platform can test and recombine them.

Google explicitly tells advertisers to upload as many **text, image, and video assets** as possible.

Media mix fragmentation keeps expanding

As budgets spread across more channels, each with different creative norms and specs, brands need more versions of campaigns to fit those environments.

AI is accelerating the expectation of output

AI lowers the cost and time required to generate variants, so leadership and platforms increasingly expect teams to produce more assets, test more combinations, and refresh creative faster. That creates its own challenges. Generative AI lets marketers scale content, but without the right structure and governance baked in, results are bound to disappoint. More on that later.

The approaching wave of LLM search

Companies are no longer creating content just for human readers or even for search engines. They are increasingly creating it to feed the LLM machine: a growing machine layer that ingests, interprets, summarizes, and surfaces information on users' behalf. What started as chatbots that can provide answers to questions can now guide consumers toward purchases.

Companies must now publish material that is clear, current, structured, and authoritative enough for AI systems to use and reuse. In practice, the rise of LLMs has turned content into both a marketing asset and a kind of machine-readable infrastructure.

This is also where the deeper challenge for CMOs rests. LLM search is the surface marketers can see today. Agentic commerce is the next one. By 2027, something else will demand its own format. CMOs are trying to address today's new challenges while also preparing for an unknown future.

A Capgemini report found 25% of consumers are already using Gen AI shopping tools in 2025 and an additional 31% have plans to do so soon.

There's a **93% increase in content demand**, according to Deloitte Digital.

It's clear that companies are under significant pressure to produce more content.

In a 2024 Capital Markets Day speech, WPP chair and former Chief Executive Officer Mark Read identified **“creative availability”** as the single biggest bottleneck in media optimization.

For CMOs, creative availability is not a production problem. It is a growth problem.

Every channel investment a CMO decides to make **depends on fresh content to perform.**

As SKUs, channels, formats, and markets proliferate, so does the content needed to convey that variety. Soon, the system collapses under its own governance weight. Lead times grow longer, revisions increase, agility erodes, and consumers stop noticing.

02



Structural Failures of the Linear Content Supply Chain

53%

In an Adobe study of 3,000 **CX executives and practitioners**, 53% said their content supply chain was still mostly linear and consumed too many resources.

47%

In addition, only 47% are using generative or agentic AI for journey design or omnichannel activation, leaving a lot of efficiency gains on the table.

Efficiency gains are still left on the table.

Human review is a legacy risk mechanism

Creatives are not historically known as strict “rule-followers.” A team of humans looking to break through a crowded marketplace can often break rules and look to create attention, sometimes by doing things with which the brand is not comfortable.

Legal, brand, and regulatory reviews exist because content is manually produced and therefore unpredictable.


Review is an insurance policy against inconsistency. Every rejected claim, reformulation, or escalation represents wasted effort. Compliance applied downstream is always expensive.

Regulatory and legal controls are applied after content is written. This guarantees delay.

However, the necessary, good intentions for content review have ballooned into enterprises of their own.

In that Adobe study, **47% of respondents** estimated that between 51 and 200 people can be involved in the creation, review, approval, and posting of a piece of content. Nearly one-fifth said it takes more than 200 people.

There is another issue. Humans are reviewing other humans. At scale, this becomes untenable.



A system that relies on review to ensure quality will never scale.


Offshoring is labour arbitrage, not productivity

Very few companies can or would even attempt to have copywriters of every native language on staff. That has led to offshored translation and transcreation. That reduces unit cost, but it does not improve throughput intelligence. They replicate the same linear process at lower wages. But lower cost does not equal higher capability. There are attendant issues as well given the distributed nature of this work.

The long-germinating integrated marketing plan finally ships to international markets after months of development. Foundational content flows through an asset management system, and then creatives get to work adapting it to their local needs, and send their drafts weeks or months later.

But external teams cannot respond to changes in direction or strategy in real time, and they are rarely given enough performance data to understand what changes need to be made. They are also fallible because they're not always able to enforce governance consistently.

The consequences compound across every cycle



Marketers across the globe end up making judgment calls on content without always having the full picture.

This is also why agency renegotiations rarely produce real savings. The cost line stays flat or rises while output quality drops, and the CMO ends up defending the budget without anything new to show for it.


Content treated as assets to be managed

Most organizations manage content as assets to be created, stored, and reused before eventually decaying. As long as content is treated as a static asset, performance will plateau.


Systems behave differently in a much more powerful and predictable way. They learn, adapt, and improve with use, whereas human judgment varies.

This is also where many martech investments have failed. The tools CMOs are asked to buy, DAMs, MRMs, and CDPs, weren't built to manage produced content. A new approach is needed.

03



A Transformation is Necessary



To meet the content needs of the AI era with the same playbook is like hiring more scribes while the printing press sits in the next room. The scribes get grumpier, the backlog grows, and everyone keeps asking why output is not improving.

Most content transformation programs have historically begun with cost pressure. Offshore the translations, commoditize the production, compress the review cycles. When the pressure eases, re-shore some of the work. The cost line moves, the operating model stays the same.

That framing has masked a more fundamental reality: the linear content supply chain was always inefficient, and now inefficiency is no longer the worst of it. The model also cannot produce content that learns. Every asset is a one-off, every cycle resets the brand to zero, and effectiveness depends on whether any individual marketer happened to make the right call this time.

Every approval step is therefore not a sign of quality. It is evidence of systemic failure upstream. Every re-shoring conversation is not a strategy. It is a stall.

What CMOs need is a system that compounds: brand rules encoded once, products represented as living data, and content that gets sharper with every market signal it absorbs. That is not a cost program. It is a performance program.

04




The System-Centric Content Supply Chain Architecture

The future content supply chain removes the individual from the center of production

Humans design the system that produces the content.

This model is non-linear, data-driven, and continuously adaptive. It must follow a certain architecture to produce value.



The platform sits on three capabilities, all governed by the same intelligence layer.

Digital shelf

Generates the complete package for PDP and ecom from a digital twin source

Adaptation


Generates the localized variants, formats, and channel cuts that scale a single creative across markets.

Originals

Generates new content from the same governed inputs when no source asset exists.

Twins are the source of truth, adaptation is the volume engine, Originals is the creative extension. The architecture on the next page is what makes the three operate as one system.

Design systems replace briefs



Most CMOs already have brand guidelines, a design ops function, and a DAM. The system-centric design system is what creates an executable, repeatable version of them.

Content creation begins with a formalized design system that encodes brand rules (why, where, when, what, what-if) at the beginning, eliminating the need for last-mile debates. The system also includes claim libraries and lexicons (governed, in-market legal content) at the root level, ensuring every piece of content adheres to specific guidelines, especially in highly regulated marketplaces.

In this model, compliance is encoded upstream rather than checked at the end. Rule-based validation, confidence thresholds, and auditability are built into the system from the start, so content is governed as it is generated, not after it has already moved through production.

It also places tone and stylistic constraints and guidelines about modular content primitives (content atomization) at the foundational level. Design becomes a one-time investment that scales infinitely.

Data feeds replace stakeholder opinion


In this new architecture, authoritative data becomes the source of truth. The system automatically draws from product and PIM data, regulatory constraints, market performance signals, and audience and channel requirements in real-time, ensuring every piece of content has the right information to become meaningful and actionable.

PRODUCT AND
PIM DATA

REGULATORY
CONSTRAINTS

MARKET
PERFORMANCE
SIGNALS

AUDIENCE
AND CHANNEL
REQUIREMENTS



Configurators replace human judgment

Decision logic containers replace interpretation. Configurators determine:

WHAT
CONTENT IS
ALLOWED
WHERE


WHICH
CLAIMS ARE
VALID PER
MARKET

HOW TONE
ADAPTS BY
AUDIENCE

WHICH
FORMATS
RENDER PER
CHANNEL

Content is generated
from the product,
not written about it.

Digital product twins replace manual authoring



In a system-centric environment, the product generates content through its digital product twin.

This canonical content engine generates content directly from the product's underlying data, rules, and attributes.

When the product changes, the product-rendered content updates automatically: product images, pack shots, retail visuals, marketplace descriptions. If a soft drink manufacturer changes ingredients or the color of the beverage, every asset generated from the twin reflects the change without human intervention.

Likewise, the system ingests regulation shifts and produces compliant language across markets instantly.

In this model, content is no longer a static output that can easily decay over time. It becomes executable logic: structured, responsive, and tightly connected to the product itself.

AI execution replaces production and translation

The right systems approach combines the generative power of LLMs with the discipline of traditional structured content systems.

AI works inside defined constraints, drawing on governed inputs, approved rules, and validated terminology to support language generation, variant creation, translation, and localization.

The system works most directly where content is generated from the product itself: product detail pages, retail media, packaging artwork, marketplace variants.

The product is the anchor, the digital twin is the source, and the output is high-volume, on-brand, and continuously updated as the product changes.



The same intelligence layer extends naturally into campaign content where the product remains the anchor: hero imagery, e-commerce campaigns, social posts, and video cutdowns where the asset is product-rendered.

And into adjacent surfaces, including influencer-style content, UGC-format outputs, and lifestyle storytelling, as the same governance (brand rules, product truth, market signal) is applied to increasingly less product-anchored content.

These are expansions from the core, not separate systems.

On their own, LLMs can be flexible but unreliable; on their own, traditional systems can be reliable but rigid.

Together, they make it possible to scale content with both speed and control.

AI executes within defined constraints:



LANGUAGE GENERATION



VARIANT CREATION



TRANSLATION AND LOCALISATION



CHANNEL FORMATTING



CONTEXT CURATION (LIFESTYLE)

Human expertise goes where it's most helpful


Human involvement shifts toward exception handling, with people stepping in when something falls outside approved parameters rather than reviewing every asset by default.

The result is that content is effectively pre-approved at the system level, **making review optional in most cases instead of mandatory every time.**

How it works in practice

This model is the one Grip has implemented for global brands including The Coca-Cola Company, LVMH, Beiersdorf, and Kellanova to adapt marketing, e-commerce, and retail visuals at scale.

Adaptation was the natural first application. Inputs, brand guidelines, and constraints all provide something concrete to work against. Original content is the harder version of the same problem, and it has been the one our customers kept asking us to solve.




We recently introduced **Grip Originals at NVIDIA GTC**, extending the same infrastructure into the generation of original content through controlled AI workflows.

The Originals solution does not replace human insight, artistic vision, or storytelling. Those stay where they have always been: with the creative leadership of the brand and its agencies. What Originals does is extend that work across markets, moments, and channels, generating new content within the same brand rules, product truth, and legal constraints that govern the rest of the system.

Generating original content puts more weight on the orchestration layer. Brand rules, legal constraints, and model selection all have to hold, sometimes without a reference image to anchor against. The intelligence layer is what makes that possible.

05



What Organizations Must Change

Technology alone is insufficient


The harder work is building the operating model that governs this entire new system.

A production system that encodes brand and compliance rules only works if the organization is willing to change how it governs, measures, and structures the teams that use it.

That starts with **governance**, which must move upstream into the design of the system itself. The rules get encoded once and enforced on every output, rather than relitigated per project.

Organisations must:

- Shift governance from approval to design
- Redefine the role of legal and compliance
- Abandon FTE-based production metrics
- Measure speed, consistency, and performance
- Re-engineer Marketing Operations and redefine Brand–Technology–Design–Operations relationships



This is an operating model transformation, not a tooling upgrade.

Legal must define the constraints the system operated within

Legal and compliance must move upstream too.

Their role is no longer to gatekeep finished work but to define the constraints the system operates within.

They become architects of the rulebook, not auditors of the output. It is a more strategic seat, but it requires giving up a familiar one.

Production metrics need the same rethink. Counting full-time equivalents made sense when humans were the throughput. In an automated pipeline, FTE counts measure the wrong thing.


What matters is how fast content reaches the market, how consistent it is across regions, and how well it performs once it gets there. Those are the numbers that tell you if the system is working.

Then there is the organization itself. Marketing Operations, Brand, Technology, and Design have historically sat in separate lanes with handoffs between them. Of course, no real CPG organization is that tidy.

Most are matrixed across regions, brands, and categories, with overlapping authority and legitimate reasons to keep these functions distinct. But the lanes are still there, and they have produced a content pipeline built on sequential custody.


Work moves between functions rather than through them, potentially eroding context at every step. An agentic content system does not tolerate handoffs well. These functions need new structural relationships, with shared ownership of the pipeline rather than sequential custody of it.

None of this is a tooling upgrade. It is an operating model transformation, and the companies that treat it that way are the ones that will see the returns the technology can actually deliver.



06

What Disappears in This Model



Unlike other “transformations” that prioritize optimization of existing services, this aims to remove several of them from the process.

The work is moving from manually producing and routing every asset to designing the rules, safeguards, and inputs that allow content to scale responsibly. This will create a shift in job responsibilities, one that many creatives will welcome. Instead of getting bogged down by laborious tasks not better handled by automation, they can focus their efforts on higher-level thinking and strategic planning.

It is a complicated and difficult process requiring hard but necessary decisions about which structures to dismantle. It involves changing the value of work that many people and teams have built careers around. Roles and services that were spokes in the global content operation, such as local translation vendors, transcreation studios, production traffic managers, and local market editing teams, are becoming less central as more of that work is absorbed into systems.

In place of labor-heavy coordination models, organizations are putting more weight on system governance, rule management, data stewardship, and exception handling.

As with other major disruptions, AI is creating discomfort about how it will change organizations’ relationships with workers. It produces a tendency to keep the status quo to avoid unpopular decisions. But it is apparent that companies should embrace the future before their competitors do.

This shift is happening; companies have an obligation to handle it with enough clarity and humanity to help people move into the work that comes next.

**DISAPPEARING ROLES
AND SERVICES**

Transcreation
studios

Production traffic
managers

Blanket legal
review cycles

Local market
editing teams

REPLACED BY

System
governance

Rule
management

Exception
handling

Data
stewardship

Automated
adaptations

07



Commercial Impact



The system-centric model fundamentally changes the structure and economics of creating content

In the traditional model, content is treated as a series of discrete deliverables, each with its own cost, timeline, approvals, and production path. It is often priced on a cost per asset basis: one landing page, one brochure, one translation, or one campaign variation.

But that logic frays when content is generated from systems rather than assembled manually. Each additional piece of content approaches a cost of zero.

At that point, the more relevant unit is no longer the asset. It is the product-market combination.

Companies will increasingly need to think about **the cost of enabling a product to go to market across regions, channels, and audiences**. The cost of each individual piece of content matters less.

As a result, variable labor costs begin to fall because fewer assets need to be touched by large chains of human producers, coordinators, and reviewers.

Time-to-market collapses from weeks or months to hours or minutes because content can be configured and generated on demand instead of routed through a long production workflow.

Grip has seen this play out with global brands managing hundreds of SKUs across dozens of markets.

Opella scaled centralized content across 260+ products and 15 markets, using automation to make content 30% more cost-efficient while giving local teams more autonomy.

THE REALITY

260+
products

15
markets

Centralized
content
demand

THE OUTCOME

30%
more
cost efficient

Local teams
gained more
autonomy

Automation
enabled scale

Governance, contrary to the fears many organizations initially have, can actually become stronger. Rules, guardrails, and approvals can be encoded into the system itself, making compliance more consistent than in fragmented manual processes. Performance compounds because the system itself learns.

Market performance signals (engagement, conversion, channel response) flow back into the configurator logic. The configurator updates the rules that govern the next cycle of outputs. Each cycle produces content that is more likely to perform in market than the one before it.


The brand owns the intelligence layer, which means the learning compounds inside the brand, not inside a vendor. Models can be swapped. Channels can shift. The intelligence layer outlasts both. The challenge, of course, is making this transition without breaking the business in the process.

The smartest enterprises do not rip out the old model overnight. They move in phases



The first step is a parallel run, where AI-native and API-first content production platforms operate alongside legacy production for selected products, categories, or markets. This is the equivalent of proving the new operating model in a controlled environment before expanding it more broadly.

In practice, this is what early adopters are already doing. Grip has run parallel pipelines alongside **legacy production** for LVMH Maisons, Mars Snacking, and others, proving the new operating model on a defined slice of the portfolio before scaling it.





From there, the organization can move toward product-led content, where the digital product twin becomes the primary source for messaging, claims, attributes, and variations.

The Coca-Cola Company offers a working example. Serving 2.2 billion drinks a day across more than 200 markets, the company has codified its **global portfolio into digital twins**, with every detail of a bottle or a cap captured as structured, governed data. Traditional photography has been replaced by a system that renders consistent, on-brand content across regions, languages, and campaigns. When the product changes, the visuals update with it. Creative changes that once required coordinated photoshoots, agencies, and translation vendors now happen inside the system. The same infrastructure is now extending into internal selling tools, inventory planning, and store simulations, which is what happens when content stops being a static asset and starts being executable logic.

The next stage is market-driven optimization, where performance data feeds back into the configurator logic, improving output systematically rather than through ad hoc revisions.

08




Legacy decommissioning

Eventually, legacy decommissioning follows. Manual production layers and offshore services are reduced or retired as the new system proves it can deliver faster, more consistently, and at lower cost. That is the uncomfortable reality of this shift, but for many enterprises, it is quickly becoming the only viable one.

Where This Goes

The future of content
is not faster production.

**It's a
system that
learns.**



That asset compounds. Volume keeps rising, surfaces keep changing, models keep changing. CMOs who keep buying point solutions pay for the same migration every eighteen months. CMOs who own the system layer get an asset that outlasts model swaps, vendor churn, and channel shifts, and produces content that performs better, not just costs less.

The work in front of the CMO is clear: design the system layer, choose the components that operate inside it, and build a content function measured on uplift, not output.