BERNSTEIN
SOCIETE GENERALE GROUP

17 September 2026

Korea Internet

Korea Internet: Sovereign AI... Sovereignty or symbolism?


Min-Joo Kang
+852 2123 2644
minjoo.kang@bernsteinsg.com


Robin Zhu
+852 2123 2659
robin.zhu@bernsteinsg.com


Charles Gou
+852 2123 2618
charles.gou@bernsteinsg.com


Makoto Morozumi
+81 3 6777 6972
makoto.morozumi@bernsteinsg.com

The goal is not to win the AI race, but to avoid dependence on its winners. Korea's sovereign AI push is increasingly framed as a race to create a domestic AI champion. We think that framing misses the central point. The real objective is not necessarily to produce a Korean equivalent of ChatGPT, Claude, or Gemini. It is to create a credible alternative that preserves strategic flexibility. In that sense, sovereign AI should be valued first as an option on operational resilience, not as evidence of platform leadership.

A fallback creates value. Owning a domestic model does not matter much if Korean consumers and companies rely on foreign agents for their most important workflows. Real sovereignty exists when Korea can fall back on a controllable alternative, maintain access to critical AI capabilities, and avoid dependence on any single provider. That alone can create meaningful value through resilience, bargaining power, and infrastructure demand, even without producing a globally dominant platform.

The Hancom playbook... strategic independence is a goal. Hancom illustrates one possible outcome rather than a definitive blueprint. It hasn't become a software leader, but it preserved a strategically important domestic alternative to Microsoft Office. Sovereign AI could follow a similar path. A domestic model may become strategically valuable in government, defense, and regulated industries while strengthening Korea's negotiating position with global AI vendors. That would still be a policy success, even if most Korean users ultimately spend their AI time on global frontier ones.

The real battle is for "Context". The more important competition is taking place at the application layer. Global AI models are turning credibility into frequency, and frequency into proprietary context, execution access, and deeper engagement. AI agents could create an even stronger flywheel than YouTube's capture of Korean user attention, because they learn not only what users consume, but how they work, decide, and delegate. Over time, the winner may not be the model with the highest benchmark score, but the one with the deepest understanding of the private user context.

Naver and Kakao: Context must become action. This is where Naver and Kakao matter. Naver owns "commercial context" through search, shopping, reviews, maps, reservations, and payments. Kakao owns "relationship and persona context" through messaging, mobility, gifting, and payments. They are potential execution moats, and orchestration layer matters here. The challenge is turning that context into completed tasks before global agents become the default starting point for Korean user intent.

For Naver, infrastructure offers the safer path as it can benefit regardless of which AI agent wins, providing a picks-and-shovels way to monetize AI adoption. But infrastructure captures demand, not necessarily the highest-value economics. The richest value still sits with whoever owns user intent, context, and action. **For Kakao**, the opportunity is more application-driven. KakaoTalk already owns frequency and relationship context, but frequency alone is not a moat in the AI era. The challenge is converting communication into delegated action. **For investors seeking a more direct sovereign AI play**, an **Upstage win** would reinforce its standing as a leading domestic AI contender while increasing the visibility of **SK Networks (001740 KS)**' investment value (c.12% stake in Upstage).

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.

First Published: 16 Sep 2026 21:00 UTC Completion Date: 16 Sep 2026 02:04 UTC

BERNSTEIN TICKER TABLE

Ticker	Rating	Cur	15 Sep 2026 Closing Price	Price Target	TTM Rel. Perf.	Adjusted EPS Cur	2025A	2026E	2027E	Adjusted P/E (x) 2025A	2026E	2027E
CPNG (Coupang)	U	USD	14.68	12.00	(70.8)%	USD	0.11	(0.54)	0.21	130.3	(26.9)	71.3
352820.KS (Hybe)	O	KRW	173,100	360,000	(61.3)%	KRW	(5,713.11)	4,637.46	9,275.75	(30.3)	37.3	18.7
035420.KS (Naver)	O	KRW	206,000	330,000	(33.6)%	KRW	13,065	13,258	15,714	15.8	15.5	13.1
035720.KS (Kakao Corp)	O	KRW	34,500	55,000	(66.4)%	KRW	1,113.49	1,267.70	1,848.11	218.4%	28.8%	32.6%
SPX			7,585.73									
ASIAX			1,916.09									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

035720.KS estimate is Reported EPS; 352820.KS valuation is Reported P/E (x); 035720.KS valuation is EPS Reported CAGR;

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We have **Outperform** rating on **Naver** and **Kakao**.

We view **Naver** as a differentiated AI infrastructure and commerce-orchestration play. Sovereign AI, cloud, and data-center investments provide exposure to AI demand regardless of which agent wins, while Naver's commercial context offers optionality at the application layer. The key debate is whether Naver can remain the owner of intent and transaction flow, rather than becoming a data supplier to global AI agents.

We view **Kakao** as an application-layer AI opportunity built on Korea's largest communication platform. KakaoTalk's distribution advantage is clear, but frequency alone is not a moat. The key debate is whether Kakao can convert relationship context into trusted AI-driven actions across commerce, mobility, payments, and daily workflows. Success would create a new monetization layer; failure would leave AI as a feature rather than a platform shift.

We believe Kakao's AI opportunity remains intact, but it is unlikely to be the primary driver of investor sentiment in the near term (till Q1 2027). Instead, investors' attention is expected to center on governance and restructuring developments. Greater clarity on these issues, alongside evidence of AI execution, could eventually allow investors to reassess the company's longer-term monetization potential.

An Upstage victory in the Korean foundation model contest would primarily benefit SK Networks (001740 KS; SK Group ICT distribution and operator of Walkerhill Hotels & Resorts) and, to a lesser extent, Kakao. SK Networks stands out as the clearest listed beneficiary given its c.12% ownership stake, which could be re-rated on the back of stronger government support, enterprise AI adoption, and increased IPO optionality. **Kakao**, with a c.8% stake, would also benefit from higher implied value for its investment while gaining strategic exposure to Korea's sovereign AI agenda. However, given the relatively small size of the stake versus Kakao's overall market value, we expect any share price impact to be secondary to the broader implications for its AI positioning. **KT (030200 KS, Korean telco operator)**'s 2.8% stake offers some exposure but is unlikely to move the needle for the investment case.

VALUATION COMPS TABLE

EXHIBIT 1: **Asia Internet & AI: valuation summary**

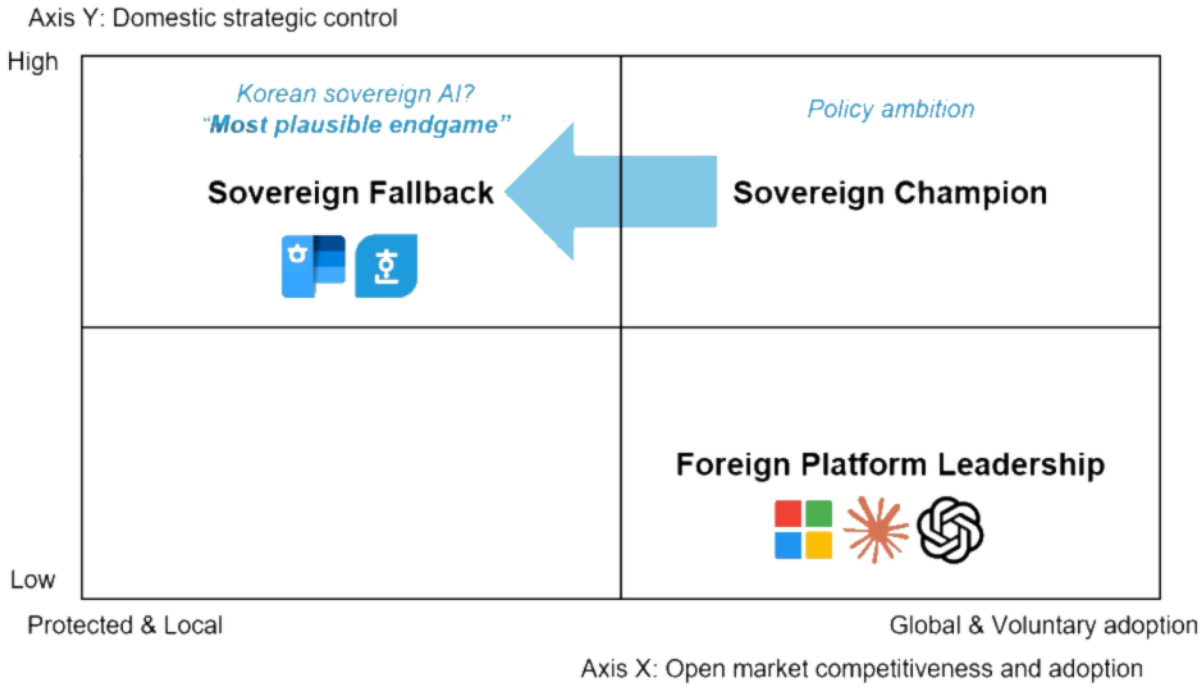
	Rating	Price target	Last price	Crncy	Market cap (US$mn)	PE			EV/sales		
						2026E	2027E	2028E	2026E	2027E	2028E
China Internet coverage											
Tencent	O	760	430.60	HKD	512,276	12.5x	11.9x	10.6x	4.4x	4.0x	3.6x
PDD	M	110	79.38	USD	112,989	8.0x	6.6x	5.7x	0.6x	0.3x	0.1x
Meituan	M	85	75.05	HKD	59,561	75.8x	11.9x	8.2x	0.7x	0.5x	0.4x
NetEase	O	160	117.99	USD	75,543	13.6x	11.7x	11.2x	3.0x	3.1x	3.2x
Boss Zhipin	O	19	15.45	USD	7,335	12.1x	10.3x	9.0x	3.0x	2.3x	1.8x
JD	O	40	27.26	USD	37,232	7.2x	6.4x	4.7x	0.2x	0.2x	0.1x
Alibaba	O	180	109.23	USD	271,505	27.4x	16.8x	11.6x	1.8x	1.7x	1.5x
Z.ai	O	1,500	721.00	HKD	41,168	-116.5x	581.8x	81.7x	33.0x	9.3x	4.9x
Minimax	M	320	252.40	HKD	10,597	-15.2x	-14.0x	-12.9x	14.4x	4.4x	2.4x
China Internet other											
Kuaishou			30.44	HKD	16,792	13.1x	12.0x	10.1x	0.6x	0.6x	0.6x
Bilibili			15.74	USD	6,590	16.1x	12.5x	9.9x	0.9x	0.8x	0.8x
TME			8.11	USD	13,475	8.8x	8.1x	7.6x	1.7x	1.6x	1.5x
Baidu			91.90	USD	31,402	16.0x	13.1x	11.1x	0.8x	0.8x	0.7x
VIPshop			12.44	USD	5,975	6.0x	5.0x	4.9x	0.2x	0.2x	0.2x
Tencent Music			8.11	USD	13,475	8.8x	8.1x	7.6x	1.7x	1.6x	1.5x
Trip.com			39.08	USD	24,609	11.4x	9.7x	8.6x	1.8x	1.6x	1.4x
KE Holdings			16.70	USD	19,184	15.4x	13.7x	12.7x	1.1x	1.0x	1.0x
Asian Internet											
Naver	O	330,000	206,500	KRW	23,200	15.6x	13.1x	11.3x	1.7x	1.2x	0.9x
Kakao	O	55,000	34,700	KRW	11,355	27.4x	18.8x	15.6x	1.0x	0.8x	0.5x
Hybe	O	360,000	175,400	KRW	5,585	37.8x	18.9x	16.2x	1.5x	1.5x	1.4x
Coupang	U	12	15.11	USD	27,163	-27.7x	73.4x	40.6x	0.6x	0.5x	0.5x
Sea Ltd.			108.68	USD	66,565	29.9x	22.3x	17.4x	2.0x	1.6x	1.4x
Grab			3.02	USD	12,319	22.7x	19.2x	13.9x	2.0x	1.7x	1.4x
US Internet											
Amazon			253.54	USD	2,734,762	17.4x	20.4x	16.2x	3.4x	3.0x	2.6x
Alphabet			349.39	USD	4,251,164	18.3x	21.8x	17.8x	9.6x	7.6x	6.2x
Meta			665.60	USD	1,695,620	19.3x	17.6x	14.3x	6.8x	5.6x	4.7x
Netflix			80.32	USD	334,448	22.5x	21.0x	17.6x	6.7x	6.0x	5.4x
Uber			72.63	USD	148,351	21.9x	16.3x	13.4x	2.7x	2.4x	2.1x
Spotify			556.31	USD	114,507	39.1x	31.2x	25.2x	4.6x	4.0x	3.5x
DoorDash			203.54	USD	88,193	37.2x	29.2x	21.1x	4.8x	4.0x	3.4x

Note: The valuation multiples of our China and Asian Internet coverage are based on Bernstein estimates; the other companies shown reflect Bloomberg consensus estimates. Close price as of September 15.
Source: Corporate reports, Bloomberg, Bernstein estimates and analysis.

DETAILS

KEY CHART

EXHIBIT 2: **The most plausible endgame for Korean sovereign AI is a Hancom-like outcome: a credible domestic alternative that delivers strategic resilience and bargaining power without becoming a global AI leader.**



Source: Corporate reports, Bernstein analysis.

KOREA INTERNET: SOVEREIGN AI... SOVEREIGNTY OR SYMBOLISM?

The real value of sovereign AI is strategic resilience

Korea can build a sovereign AI option without building a sovereign AI winner. This is the central distinction. Owning a domestic foundation model does not, by itself, constitute sovereignty. Sovereignty exists only when Korea can credibly rely on a domestic or otherwise controllable alternative to preserve critical operations, reduce dependence on foreign suppliers, and negotiate better terms on access, pricing, and data governance. Korean policy thinking is beginning to move in this direction, away from simple model ownership and toward access sovereignty and operational sovereignty: the ability to retain access to essential AI capabilities and keep critical functions running when a particular model or cloud provider becomes unavailable.

This makes Korea's sovereign AI push neither meaningless nor sufficient. At its best, the program can create national insurance against supply disruption, strengthen Korea's bargaining position with global AI vendors, and support domestic capability in strategically sensitive sectors. At its worst, it can produce expensive models that survive only through public procurement while Korean consumers and enterprises delegate their most valuable work to foreign agents. The bear case is therefore not that Korea builds nothing. It is that Korea becomes sovereign in infrastructure, symbolic in applications, and dependent at the interface layer.

The distinction matters for investors because each outcome has a different economic value. Strategic resilience can justify government spending and support domestic infrastructure revenue without creating a global platform. Protected B2G adoption can sustain local vendors but does not prove voluntary product-market fit. The greatest application-layer upside requires something harder: users must repeatedly choose a Korean agent, entrust it with proprietary context, and allow it to execute meaningful tasks. Until that happens, the program should be valued as an “option on resilience” rather than evidence of a new consumer franchise.

Without substitutability, sovereignty is symbolic

The value of sovereign AI lies less in beating the frontier than in creating a credible outside option. Korea's policy debate has focused heavily on ownership: domestic foundation models, domestic compute, and domestic data centers. The government-backed indigenous model program is designed to produce one or two models approaching roughly 95% of leading global performance by 2027. That target is politically legible and technically measurable, but it is not a complete measure of sovereignty.

AI is a stack, not a model. Reliable deployment depends on accelerators, cloud infrastructure, model access, data, security, workflow integration, distribution, and permission to act. A Korean model running on foreign hardware and software dependencies, unable to replace foreign services in production, may improve domestic capability without materially changing the country's strategic position. Conversely, Korea does not need to manufacture every layer locally to possess meaningful sovereignty. It needs enough substitutability across the stack to prevent any single external supplier from becoming indispensable.

The correct test is therefore credible substitutability. If a foreign vendor raises prices, restricts model access, changes data terms, or becomes unavailable for geopolitical reasons, can Korean institutions move critical workloads to another system without intolerable disruption? If the answer is yes, domestic capability creates real bargaining power. If the answer is no, the local model may still support research and industrial development, but its sovereignty value is largely rhetorical. Korea's strength in memory semiconductors and advanced manufacturing provides leverage within the US-led AI supply chain, but hardware relevance does not automatically translate into control over the application layer.

This framing also avoids the false choice between full technological autarky and complete dependence. Korea need not recreate Nvidia, hyperscalers, and every frontier lab at once. It needs a portfolio of domestic capacity, open models, trusted foreign partnerships, and migration-ready infrastructure that preserves freedom of action. In defense, government, healthcare, finance, and other regulated sectors, a locally controllable model may remain strategically valuable even if it trails the frontier. These users care about service continuity, security, auditability, and jurisdiction as much as raw intelligence.

The policy scorecard should therefore move beyond benchmark scores, GPU counts, and the number of selecte
champions." The more relevant questions are how much critical workload can shift to a controllable stack, how l
migration would take, what performance and cost penalties would arise, and whether the presence of a domesti
improved actual contract terms. A model that performs well in a government evaluation but cannot carry produc

人在草木间

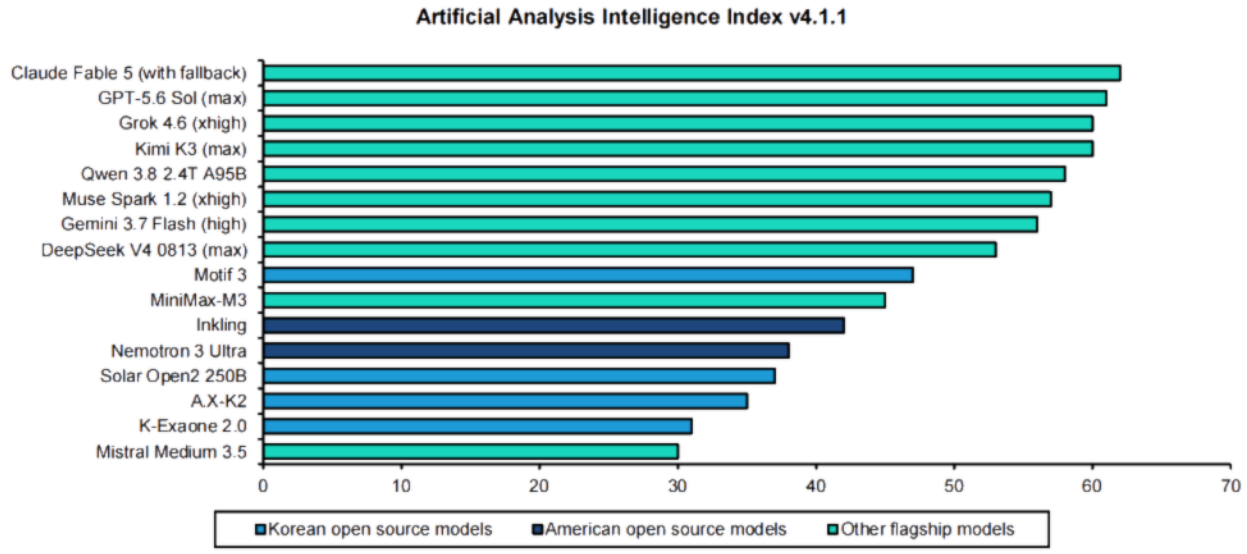
is not bargaining power. It is bargaining theater.

EXHIBIT 3: **Korea's Sovereign AI push signals a shift from model leadership to strategic resilience.**

Company	Model	Active/Total parameters	Attention type
EXAONE	K-EXAONE 2.0 750B-A37B	37B/750B	Hybrid sliding-window + global GQA
A.X K2	A.X K2	33B/688B	MLA + top-k sparse attention
upstage	Solar Open 2 250B-A15B	15B/250B	Hybrid linear (KDA) + softmax GQA
MOTIF	Motif-3	13B/314B	Grouped differential latent attention

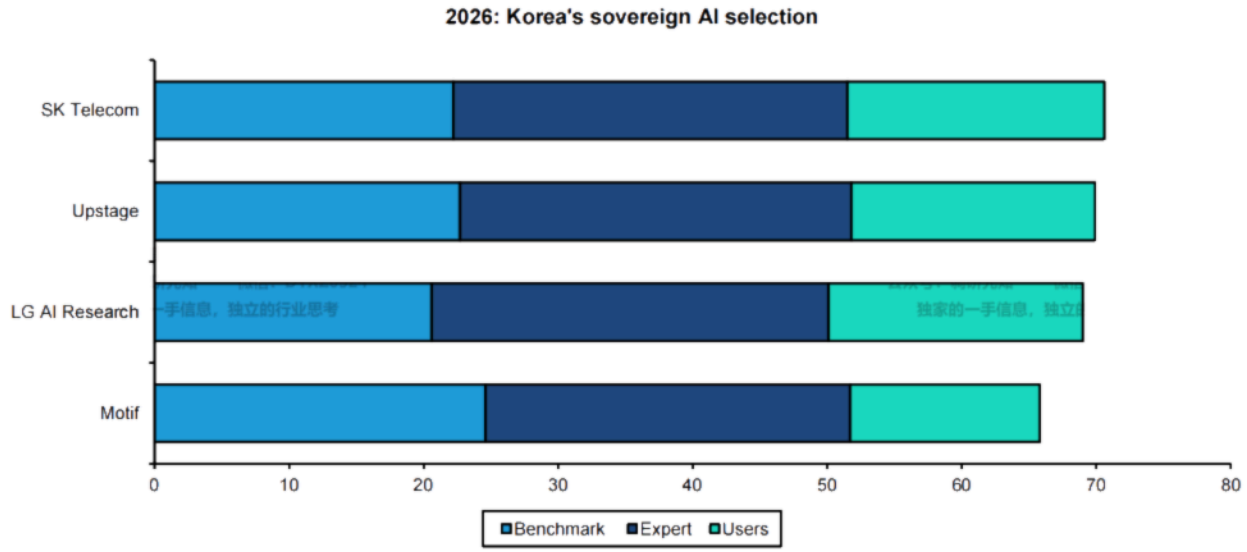
Source: Corporate reports, Semi analysis. Bernstein analysis.

EXHIBIT 4: **Motif Technologies, a relatively small startup, matched the benchmark performance of much larger global incumbents, reinforcing the view that Korea already has a viable sovereign AI option.**



Source: Artificial analysis, Semi analysis, Bernstein analysis.

EXHIBIT 5: **Motif's exclusion highlights a broader question in Korea's sovereign AI push: is technological capability alone sufficient, or do deployment readiness, user adoption, and ecosystem reach matter just as much?**



Source: Semi analysis, Bernstein analysis.

The endgame may be "Hancom"

Hancom shows that industrial policy can preserve a strategically useful domestic alternative without producing a global platform winner. In 1998, Hancom was under severe financial pressure and considered accepting a Microsoft investment tied to discontinuing further development of Hangul. A domestic rescue campaign raised capital, halted the Microsoft talks, and enabled the company to continue developing its Korean word processor. Government and public-sector purchasing subsequently provided an important demand base, while stronger enforcement against software piracy supported the market for licensed domestic software.

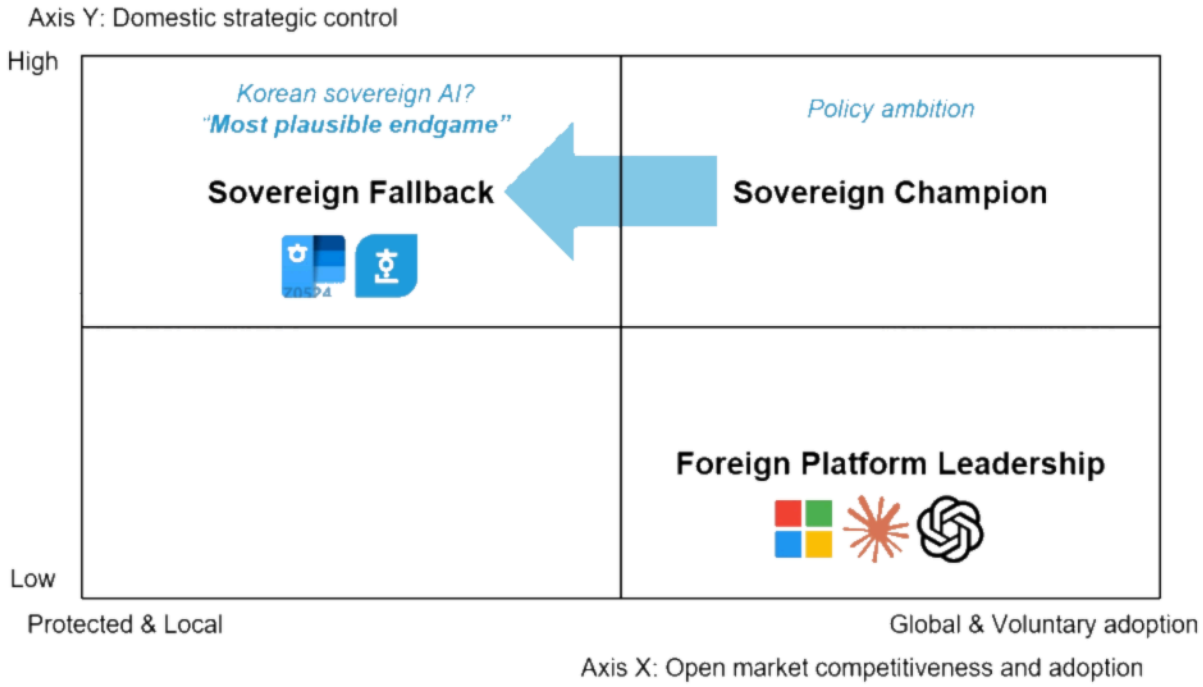
This was not an empty political gesture. Hangul preserved Korean-language software capability, public-document conventions, and compatibility with domestic administrative workflows. It prevented Korea from becoming completely dependent on a single foreign office-software provider. For government agencies and users tied to the HWP format, the product provided operational continuity and a degree of supplier choice. In that narrow but important sense, industrial policy created "genuine resilience".

The commercial ceiling was equally clear. As Korean companies became more integrated with overseas customers and partners, Microsoft formats became increasingly important, and Hancom remained much stronger in public-sector workflows than in globally connected enterprise use. The company and product survived, but the ecosystem did not become a globally dominant productivity platform. Public procurement secured a domestic enclave; it did not overturn Microsoft's global network effects.

That is precisely why Hancom is the right analogue for sovereign AI. A Korean model could become strategically important in government, defense, and regulated industries, remain commercially viable through local demand, and give the country a fallback option during a supply shock. None of those achievements would prove that the model can win open consumer or enterprise markets against ChatGPT, Claude, or Gemini. Survival, resilience, and bargaining power are legitimate policy outcomes, but they should not be confused with application-layer leadership.

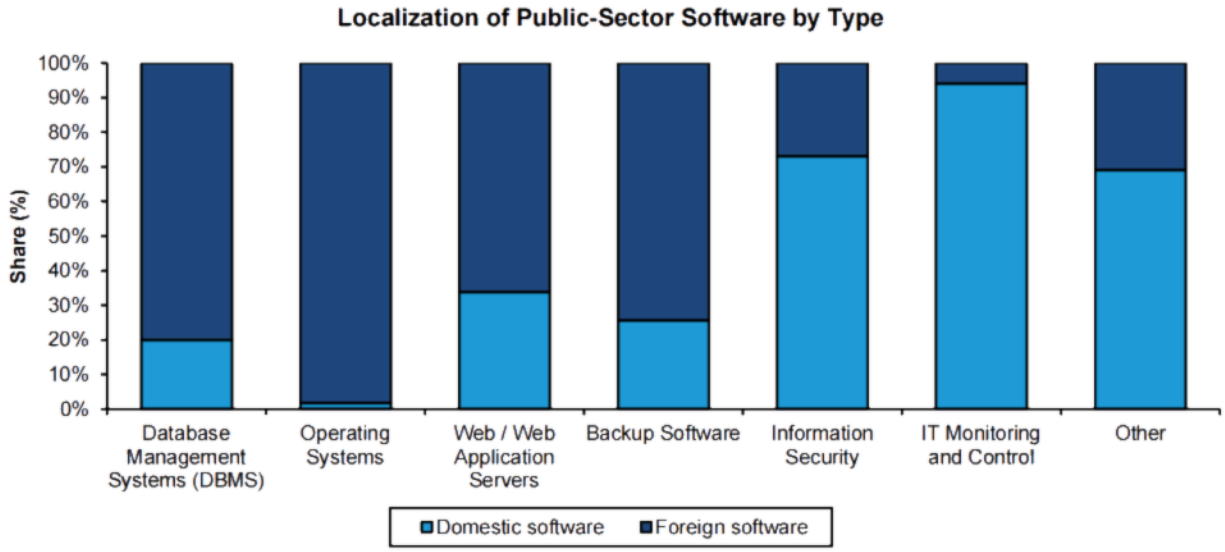
The most plausible success case may therefore be "Hancom for AI": a durable and useful domestic alternative that protects selected workflows and limits dependence, but remains commercially concentrated in protected or locally specific markets. That outcome would be more substantive than symbolism, yet far narrower than the rhetoric of national AI leadership implies. The mistake would be to call it a failure because it does not conquer the world, or to call it a global victory because public institutions continue to buy it. **Hancom achieved sovereignty as resilience, not sovereignty as supremacy. Korean AI may do the same.**

EXHIBIT 6: **The most plausible endgame for Korean sovereign AI is a Hancom-like outcome: a credible domestic alternative that delivers strategic resilience and bargaining power without becoming a global AI leader.**



Source: Corporate reports, Bernstein analysis.

EXHIBIT 7: **Korea's public sector has achieved selective software sovereignty, not stack-wide technological independence.**



Source: Korea Ministry of the Interior and Safety (2024), Bernstein analysis.

Four programs, four definitions of sovereignty

Korea's sovereign AI agenda currently spans four major initiatives: foundation models, consumer AI agents, cybersecurity AI, and infrastructure. Each pursues a different objective and defines sovereignty differently.

The Independent AI Foundation Model project, commonly known as Dokpamo, focuses on model autonomy. This is the purest expression of AI sovereignty, where success is measured by the ability to develop and operate domestic foundation models with minimal reliance on foreign technologies. The debate surrounding Dokpamo has highlighted how difficult this goal is in practice. Questions around benchmark performance, open-source components, evaluation criteria, and what constitutes genuine independence illustrate that technological sovereignty is not always easy to define, even before it can be achieved.

Modu-AI takes a very different approach. Rather than focusing on model ownership, it focuses on user adoption and distribution. The objective is to build domestic AI agents that become part of everyday consumer activity, from reservations and commerce to public services and routine task completion. In this framework, sovereignty is defined less by control over models and more by control over the interface through which users express intent and delegate tasks. The strategic objective is to ensure that domestic platforms remain relevant even if frontier AI models become increasingly globalized.

K-Mithos represents a third interpretation of sovereignty, centered on cybersecurity and critical infrastructure. Sensitive operational data, cyber-defense systems, and national infrastructure create areas where domestic control carries particular strategic value. In these environments, reliability, auditability, operational continuity, and security may matter more than possessing the most advanced frontier model. The program therefore focuses on building domestic capabilities around security use cases where dependence on external providers may become a strategic risk.

Finally, Korea's large-scale AI infrastructure initiatives focus on compute capacity, data residency, cloud infrastructure, and deployment. These projects are designed to ensure sufficient domestic capacity to support future AI development and adoption. Infrastructure sovereignty increases resilience by providing local capacity and reducing operational bottlenecks, even if it does not eliminate reliance on the global technology ecosystem.

An Upstage victory in Korean foundation model contest would primarily benefit SK Networks (001740 KS; SK Group ICT distribution and operator of Walkerhill Hotels & Resorts**) and, to a lesser extent, Kakao. SK Networks** stands out as the clearest listed beneficiary given its c.12% ownership stake, which could be re-rated on the back of stronger government support, enterprise AI adoption, and increased IPO optionality. **Kakao**, with a c.8% stake, would also benefit from higher implied value for its investment while gaining strategic exposure to Korea's sovereign AI agenda. However, given the relatively small size of the stake versus Kakao's overall market value, we expect any share price impact to be secondary to the broader implications for its AI positioning. **KT (030200 KS, Korean telco operator)**'s 2.8% stake offers some exposure but is unlikely to move the needle for the investment case.

In contrast, a Motif victory would offer no direct listed proxy. Both Motif and its parent, Moreh, remain private, while KT and AMD are only indirect investors in Moreh with undisclosed ownership stakes. Furthermore, we view a Motif win as a low-probability outcome following the dismissal of its appeal and management's earlier indication that it would not participate in the next round.

EXHIBIT 8: **Korea's sovereign AI agenda currently spans four major initiatives: foundation models, consumer AI agents, cybersecurity AI, and infrastructure. Each pursues a different objective and defines sovereignty differently.**

Program	Sovereignty layer	Principal winners	What Korea gains	What remains foreign-dependent
Dokpamo	Foundation models	SKT, Upstage, LG AI Research	Domestic model capability and fallback options	GPUs, development tools, parts of the software stack
Modu-AI	Consumer distribution	Kakao, SKT, KT	User access, frequency, transaction pathways	Compute and potentially underlying models
K-Mithos	Cybersecurity data and deployment	Naver Cloud-led consortium	Sensitive data control, vertical models, operational resilience	Nvidia compute and related tooling
AI data centers	Physical infrastructure	SK, GS, Naver	Domestic capacity, jurisdiction, data residency	Accelerators, CUDA and portions of the global supply chain

Source: Corporate reports, Bernstein analysis.

THE "FOREIGN INVASION" HAS BEGUN

And they are capturing the user credibility, frequency and user engagement.

Global LLMs are turning credibility into frequency, and frequency into proprietary context, execution access, and deeper engagement. The risk to Korean platforms is not simply that foreign AI services capture another category of screen time. AI agents could sit one layer deeper than traditional internet platforms, mediating what users ask, which sources they trust, how they work, and eventually which transactions they authorize.

YouTube provides a useful precedent because its eventual dominance was far from obvious. More viewers generated richer behavioral signals; better recommendations increased viewing time; greater reach and monetization attracted more creators; and the expanding content supply drove engagement still higher. YouTube says its recommendation system incorporates signals such as clicks, watch time, survey responses, shares, likes, and dislikes, and that recommendations generate more viewing than subscriptions or search. What began as a destination for online video became a compounding system for allocating consumer attention. The outcome in Korea has been striking. Based on our January-August 2026 dataset, YouTube accounted for 85% of time spent in the mobile entertainment vertical. YouTube did not merely defeat domestic video competitors; it rewrote the allocation of Korea's digital day.

AI agents could produce an even more powerful flywheel.

YouTube learns what users want to watch; an agent can learn how users think and work: which files matter, which sources they trust, how they want output formatted, which shortcuts they accept, and which errors they cannot tolerate. Once users begin delegating research, coding, writing, scheduling, purchasing, and administrative work, every correction, approval, tool call, and completed task can enrich personalization and improve future outcomes, depending on user consent and the provider's data architecture.

The progression could look like this: Credibility > Frequency > Context > Better Outcomes > Credibility > Delegation. Credibility comes first because users will not delegate consequential work to an agent they do not trust. Once credibility drives repeated usage, the agent begins accumulating proprietary context. Better context improves task completion, which strengthens trust and encourages users to grant broader permissions. The resulting flywheel is not merely an engagement loop; it is a delegation loop.

This also creates potentially stronger switching costs than those supporting YouTube. Leaving a video platform means losing recommendations, subscriptions, and access to a creator ecosystem. Leaving a mature AI agent may mean abandoning accumulated memory, workflow conventions, trusted-source lists, connected tools, permissions, and a shared history of what "good" looks like. The user is no longer switching software. The user is replacing a trained digital colleague.

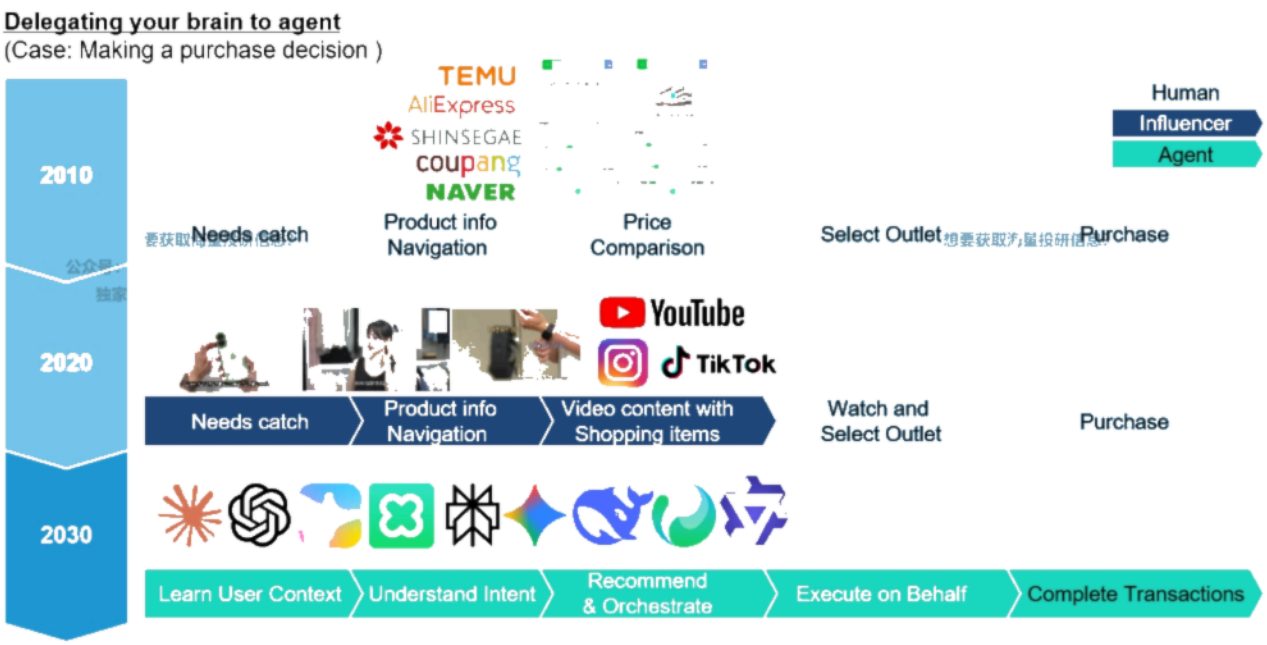
Korean adoption is already moving beyond experimentation. ChatGPT reached 17.1 million mobile MAUs in August 2026, while average monthly usage per user increased from 115 minutes to 141 minutes year over year. OpenAI also reported that its Korean ChatGPT Enterprise user base had increased 28-fold over the previous year, while weekly active users of ChatGPT Work and Codex had risen 14-fold since March. The company did not disclose absolute Enterprise user numbers, so the 28-fold increase should be interpreted as evidence of rapid growth rather than proof of market scale. Still, the direction is clear: global LLMs are extending their reach from casual queries into high-value corporate workflows.

None of this means that the first entrant automatically wins. AI is not a simple network-effect business in which every incremental user directly improves the product for everyone else. Data rights, user consent, feedback quality, memory architecture, and post-training capabilities determine whether usage can be converted into better outcomes. Users can also multi-home across ChatGPT, Claude, and Gemini, while open standards and portable memory could reduce switching costs.

For agents, early traffic is not a moat, compounded trust is.

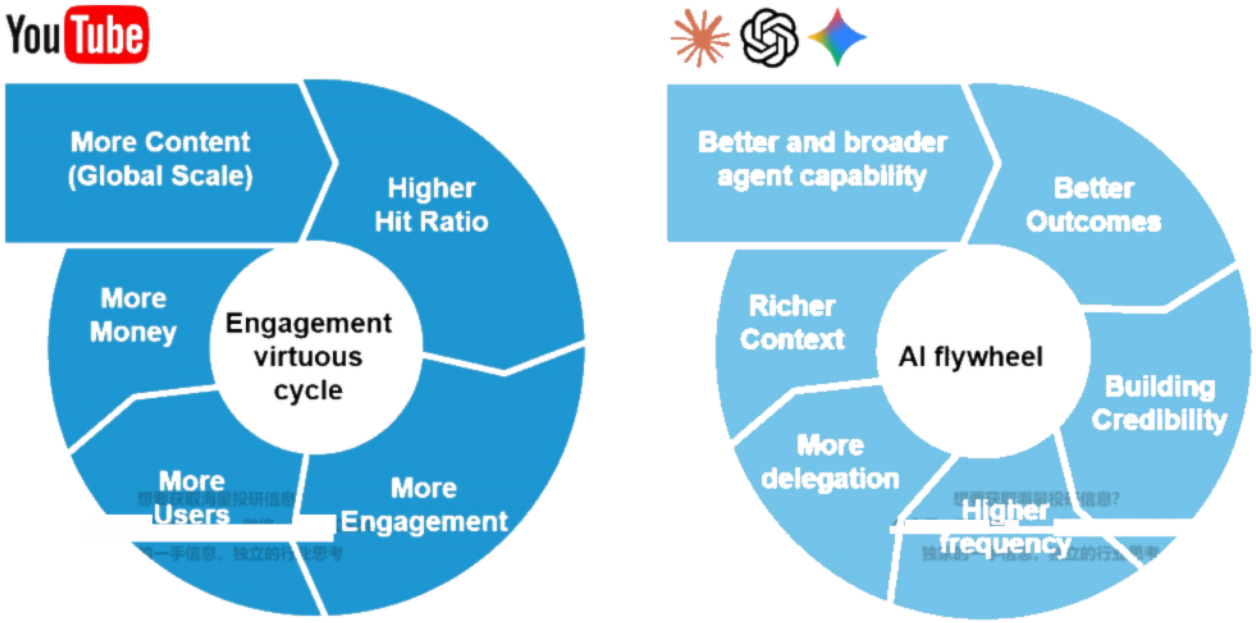
It is to become credible first in high-value tasks, earn repeated use, accumulate proprietary context, connect to essential tools, and secure permission to act. Once an agent becomes the default place where users begin their work, its advantage extends well beyond model quality. By the time agentic behavior becomes visible as a fully formed mass-market habit, the underlying contest for context may already have tilted. The first mover wins only when early usage compounds into proprietary context, execution access, and trust. Korea's AI habits are not yet settled, but the market is no longer neutral.

EXHIBIT 9: **Delegating your brain to an agent... The control point of Korean e-commerce has shifted from human discovery to influencer-driven recommendations, and may ultimately move to AI agents that learn user intent and execute purchases on their behalf.**



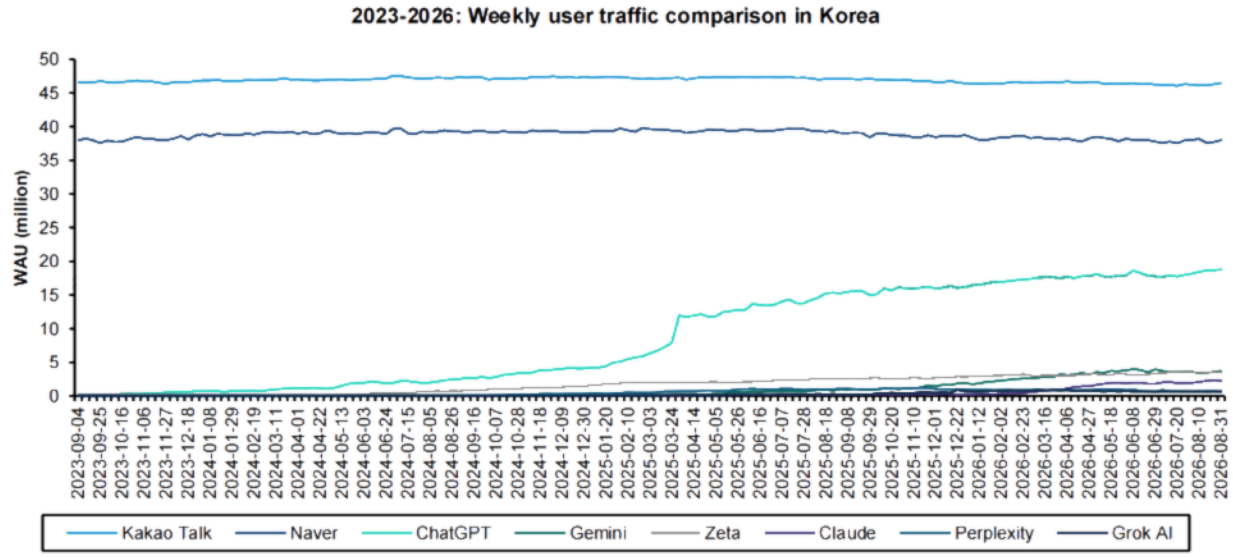
Source: Corporate reports, Bernstein analysis.

EXHIBIT 10: **Early AI usage becomes a moat only when it generates persistent context, better outcomes, and permission to act - enabled by user consent, usable data, memory, and workflow access.**



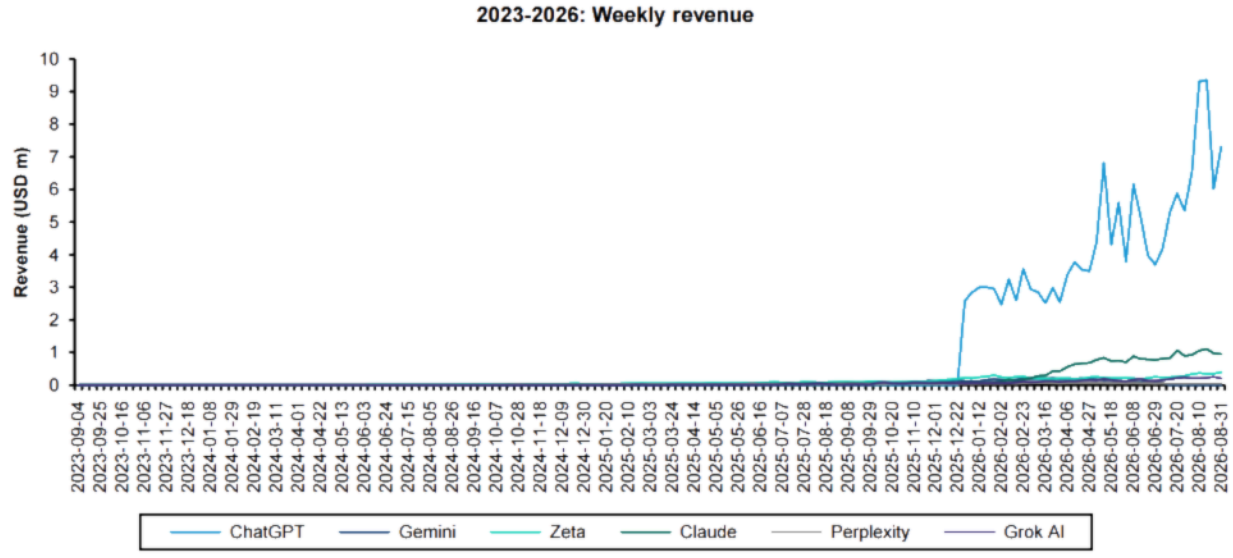
Source: Corporate reports, Bernstein analysis.

EXHIBIT 11: **ChatGPT's growing attention share in Korea suggests it is taking an early lead in the race to establish AI-native user habits.**



Source: Wiseapp, Bernstein analysis.

EXHIBIT 12: **ChatGPT and Claude are increasingly demonstrating AI monetization potential in Korea, with a combined paying-user ratio now estimated in the high single digits.**



Source: Sensor Tower, Bernstein analysis.

Local context is the last-mile moat for Naver and Kakao

As frontier AI models increasingly capture user credibility and build frequency in Korea, the key question is how local internet platforms remain relevant in the agent era. Naver and Kakao do not need to outperform global labs in general reasoning. They need to turn their unique local context into superior execution before global agents become the default interface for Korean user intent. Both companies still possess assets that are difficult for foreign model providers to replicate. Naver owns deep commercial context across search, shopping, product catalogs, reviews, maps, reservations, and payments. Kakao owns relationship context through messaging, group interactions, gifting, mobility, and payments.

Their strategies, however, are diverging. Naver is pursuing a vertically integrated stack spanning AI models, cloud, search, and commerce. Its key advantage lies not in the model itself, but in the structure of its commercial data. Product attributes, merchant information, reviews, inventory, reservations, and payments can all be linked into a transaction graph. A frontier model may generate the better answer, but Naver can still deliver the better outcome if its agent understands local inventory, interprets Korean reviews, verifies availability, and completes the transaction. Its moat is the closed loop from intent to execution.

The risk is that Naver loses control of the top of this loop. We believe this is one reason the company remains highly defensive around commerce share. If Korean consumers increasingly begin shopping, travel, or local-services queries through ChatGPT or other global agents, those platforms would own user intent, decision-making, and the feedback signals generated throughout the journey. Naver may retain valuable merchant data and transaction infrastructure, but gradually lose ownership of the user relationship. In that scenario, it would continue to participate economically, but increasingly as an execution layer rather than the primary interface.

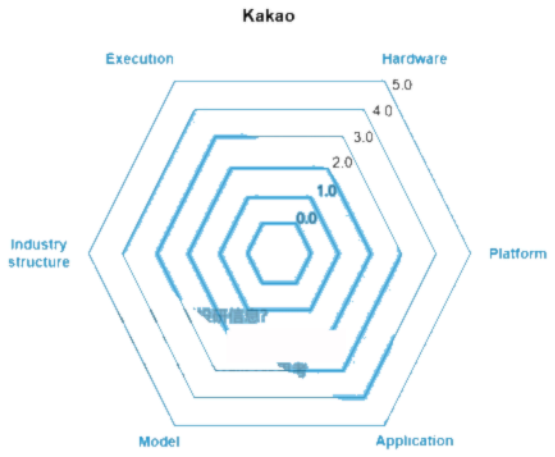
Kakao is taking a different approach, positioning KakaoTalk as an AI gateway through which multiple models and services can reach users. Its advantage lies in frequency and relationship context: understanding family ties, social networks, work relationships, and the intentions that emerge in private and group conversations. In theory, an agent could detect a dinner plan, coordinate preferences, book a restaurant, split the bill, and arrange transportation without users leaving the chat. Kanana's vision of an AI companion embedded within personal and group conversations is strategically aligned with this opportunity.

However, frequency alone does not create value in the AI era. Notification checks and short messaging interactions generate engagement, but not necessarily high-intent actions. Kakao's challenge is to convert communication into delegation across commerce, reservations, gifting, mobility, food delivery, and customer service. Expanding its agent ecosystem and integrating more leading vertical platforms may prove more important than developing the strongest proprietary model.

The real battle is turning context into action.

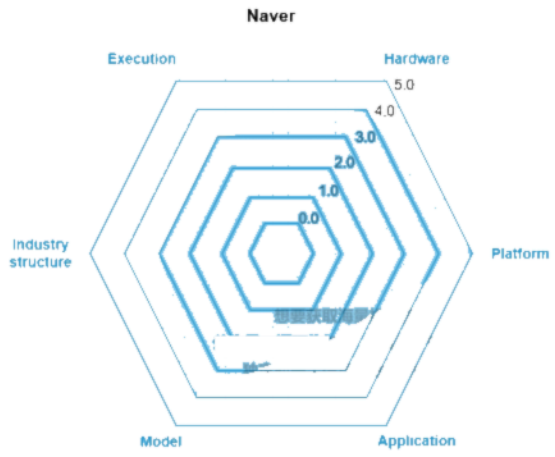
The strategic divide is therefore clear. Naver owns commercial context, while Kakao owns relationship context. Naver must prevent global agents from disintermediating its search and commerce interfaces. Kakao must turn communication into action before its frequency advantage becomes economically less relevant. Neither company needs to own the best frontier model. What matters is maintaining enough orchestration, permissions, and transaction access to remain embedded in the user's workflow when AI agents become the primary interface.

EXHIBIT 13: **Kakao's strongest AI asset is relationship and persona context, but long-term value will depend on its ability to translate that context into execution.**



Source: Bernstein estimates and analysis.

EXHIBIT 14: **Naver owns the rails of AI through platforms, models, and infrastructure; the remaining challenge is turning commercial context into completed actions.**



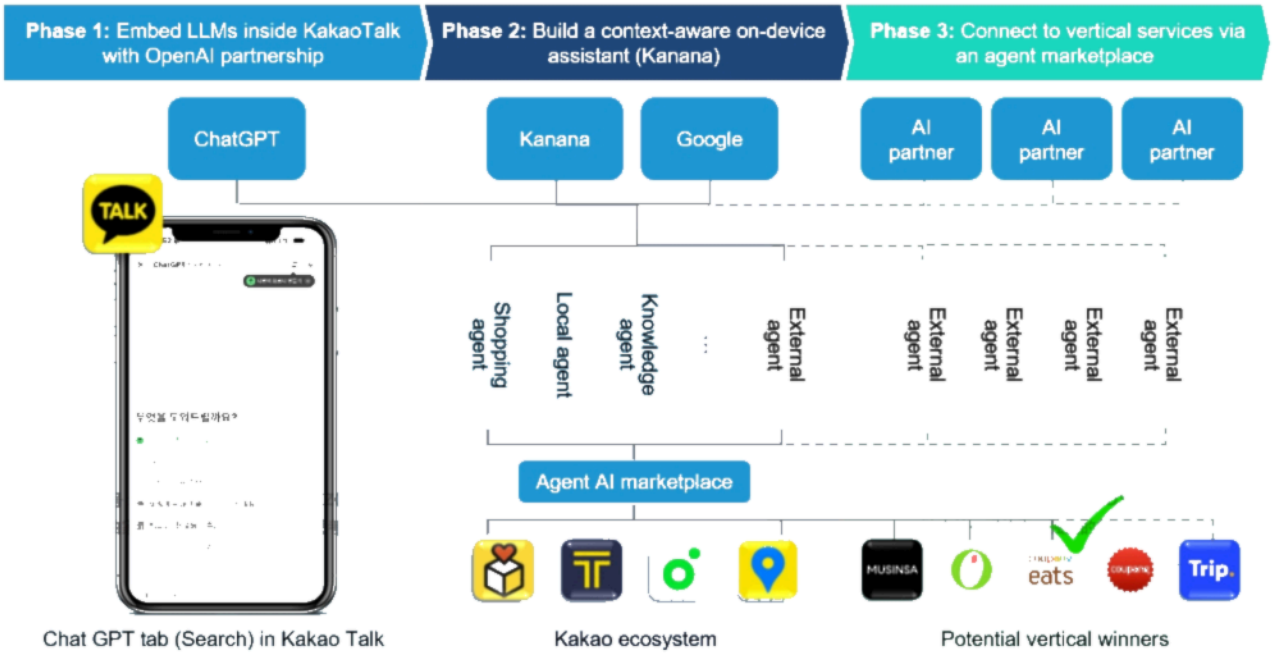
Source: Bernstein estimates and analysis.

EXHIBIT 15: **The challenge for both Korean platforms is not to build a better frontier model than global competitors, but to convert their proprietary local context into completed actions before global AI agents become the default interface for user intent.**

	Tencent ecosystem	Naver ecosystem	Kakao ecosystem
Main domain	**WeChat (Messenger)**	**Naver Search (Search portal)**	**KakaoTalk (Messenger)**
Content	▪ WeChat Mini Programs for UGC ▪ WeChat Channel for Shot form videos	▪ Naver Clip: short-form video feeds ▪ Naver Store (3P e-Commerce ecosystem) ▪ Naver News	▪ Kakao Shot form: short-form video feeds ▪ Kakao Shopping and gifting commerce
AI	▪ HunYuan LLM ▪ Yuanbao ▪ DeepSeek integration	▪ HyperCLOVA X	▪ Kanana AI Mate (On-device AI strategy) ▪ OpenAI ChatGPT
Cloud	▪ Tencent Cloud (hyperscaler)	▪ Naver Cloud (hyperscaler)	▪ Kakao Enterprise Cloud ▪ Partnerships with Google Cloud/OpenAI
Entertainment	▪ Gaming ▪ Long-form video	▪ Naver Webtoon ▪ CHZZK: live streaming service	▪ Kakao Page (Web comic) ▪ SM Entertainment (K-pop) ▪ Melon (Music streaming MAU share: 8%)
Commerce	▪ Mini Program (can access all China e-Commerce platforms) ▪ WeChat gifting commerce	▪ Naver Smart Store (Korea e-Commerce GMV share: 25%)	▪ Gifting commerce (Korea e-Comm GMV share: 5%)
Fintech	▪ WeChat Pay (GTV share: 40%)	▪ Naver Pay (6.5M MAU, c.40% of GTV share) ▪ Partnership with Dunamu	▪ KakaoPay (11.8M MAU, c.30% of GTV share) ▪ Kakao Bank (18.3M MAU)
Search	▪ Weixin Search (Yuanbao) ▪ Agent AI to be implemented	▪ Naver Portal Search (AI Briefing/Cue)	▪ Daum Search ▪ ChatGPT Integration on KakaoTalk

Note: MAU as of June 2026.
Source: Wiseapp, Bernstein analysis.

EXHIBIT 16: **Kakao's AI roadmap shifts from chat to action, from embedding LLMs in KakaoTalk to orchestrating a marketplace of agents and vertical services.**



Source: Corporate reports, Bernstein analysis.

Naver's safe bet, owning the rails... commerce agents for upside optionality

Naver's infrastructure strategy is the safer, simple and straight forward way to monetize AI demand, but infrastructure captures demand, not necessarily the highest-value economics. Data-center capacity, sovereign cloud, security, and local deployment can benefit from AI adoption regardless of which model or consumer agent wins. Government bodies and regulated enterprises will continue to require domestic data storage, auditability, and operational continuity, while Korea's broader sovereign AI strategy is expanding beyond model ownership toward semiconductors, AI factories, and deployment infrastructure.

This is the picks-and-shovels case. If ChatGPT, Claude, domestic models, and enterprise agents all drive higher inference and storage demand, Naver can participate without correctly forecasting the application-layer winner. Its local regulatory knowledge, cloud capabilities, and ability to bundle infrastructure with deployment make it a credible partner for Korean enterprises and overseas governments pursuing their own sovereign AI programs.

The strategy is safer, but not risk-free. AI data centers require substantial capital (even though the first phase of Naver Neocloud buildout can be asset light), and returns depend on utilization, power availability, customer commitments, financing costs, and the pace of industry capacity additions. Compute can become more standardized, particularly if inference efficiency improves faster than expected or competing capacity comes online. Infrastructure offers visibility, but it does not guarantee scarcity economics.

More importantly, infrastructure alone does not capture the richest learning loop. The data-center operator sees workload demand. The agent owner sees user intent, private context, corrections, permissions, and outcomes. Those signals can create personalization, switching costs, and platform control that infrastructure does not possess. If Naver emphasizes the safer layer while ChatGPT and Claude own the user relationship above it, Naver may become the landlord of an AI economy whose most attractive economics accrue elsewhere.

The optimal strategy is therefore a barbell rather than a retreat. Infrastructure provides downside protection through compute, storage, security, and sovereign deployment demand. Commerce agents provide upside optionality through control of intent, context, and transaction completion. The infrastructure business can finance the transition and reduce dependence on an uncertain consumer AI outcome; the application layer preserves the possibility of platform economics.

Naver can own the rails of Korean AI. But if foreign agents own the passenger, the destination, and the ticket, the rails may not capture the largest share of value.

EXHIBIT 17: **Naver's infrastructure strategy is the safer, simple and straight forward way to monetize AI demand, but infrastructure captures demand, not necessarily the highest-value economics.**

	Phase 1: 200MW lease	Phase 2: 200MW owned	Phase 3: 600MW hybrid
Capacity & ownership	200MW (Existing leased sites)	200MW (Owned, Sejong expansion)	600MW (lease + owned/greenfield mix)
Asset ownership	DC leased GPUs owned by SPV	Owned DC, GPUs owned by SPV, Naver owns OperatingCo SPV	Mix of leased, owned domestic, and owned foreign SPVs
Naver equity role	Roughly 1 B USD in SPV	Additional equity in Operating Co; domestic focus	Smaller per-project equity; PropCo (infra SPV) + OperatingCo (Naver 51%)
Funding mix	c.10B USD capex: 1 B NVIDIA, 9B project debt.	Sovereign/policy loans + project bonds + equity	Project finance, local partners, policy support, export-oriented capital
Contract model	Single anchor customer, 200MW take-or-pay	Additional capacity can be allocated to the anchor client	Portfolio of sovereign AI, hyper-scalers, regulated enterprises
Balance sheet impact	Most debt ring-fenced; expense flows through P&L	Still ring-fenced, but more owned assets; policy lowers cost of capital	Portfolio of off-BS SPVs, with Naver consolidating mainly operating entities
Strategic role	Prove economics	Establish Naver as domestic sovereign AI operator	Build regional AI platform and export AI DC solutions

Source: Bernstein analysis.