QUANTINGENT

Analytical Intelligence Platform

FOUNDATIONAL REPORT

QFR-001

The Chokepoint Architecture

How Semiconductor Export Controls Are Reshaping the Balance of Technological Power

Pillar Coverage	Technology · National Security · Geopolitics
QAF Version	v1.0
Classification	Institutional Distribution
Date	26 – 06 - 2026
Prepared by	QUANTINGENT Research

Executive Summary

The United States and its allies have constructed what may be the most consequential industrial policy intervention of the twenty-first century: a layered export control architecture designed to sever China's access to advanced semiconductor technology. Since October 2022, and with accelerating intensity through 2024, the Bureau of Industry and Security has expanded restrictions covering not only finished chips but the equipment, software, and human capital required to produce them. The Netherlands and Japan have followed, placing the entire frontier of chipmaking — concentrated in three firms across three countries — effectively under Western strategic control.

This report applies the Quantingent Analytical Framework (QAF) to examine the chokepoint architecture across all three pillars. The technology analysis maps where genuine vulnerabilities exist in China's semiconductor supply chain, and where the narrative of effective containment is overstated. The national security analysis interrogates the dual-use implications of advanced AI hardware and the military capability feedback loop that drives U.S. urgency. The geopolitical analysis examines the trilateral alignment that makes this regime function — and the fault lines that could fracture it.

The central QAF question — does the export control architecture accelerate Chinese self-sufficiency or effectively delay it? — does not have a clean answer. The evidence points to a more uncomfortable conclusion: the controls are working on the timelines that matter for near-term military balance, and failing on the timelines that matter for long-term economic competition.

> *The chokepoint architecture works precisely because semiconductor manufacturing is the most geographically concentrated critical technology in human history. Three firms — TSMC, Samsung, and Intel — produce virtually all advanced logic. One firm — ASML — holds a monopoly on the equipment required to make them.*

This report is structured in five sections: technology signal and capability assessment; national security doctrine and dual-use dimensions; geopolitical power distribution and alliance architecture; cross-pillar intersection analysis; and three divergent scenarios for the period through 2030.

Pillar I — Technology

1.1 Signal Layer: The Frontier and the Gap

Advanced semiconductor manufacturing operates at technology nodes measured in nanometers. The current frontier sits at 3nm and below, with TSMC producing Apple's M-series and Nvidia's Blackwell architecture at this node, and 2nm processes scheduled for volume production in 2025. China's domestic industry, anchored by SMIC, achieved a significant milestone in 2023 by producing 7nm chips — the Kirin 9000S found in Huawei's Mate 60 Pro — using workarounds to equipment restrictions. This was a genuine technical achievement and correctly generated concern in Washington.

It should not, however, be read as evidence that containment is failing. The relevant metric is not whether China can produce 7nm chips. It is whether China can produce them at yield, at volume, and at a cost that makes them competitive. Informed assessments suggest that SMIC's 7nm process suffers from significantly lower yields than equivalent processes at TSMC or Samsung, requires far more manufacturing steps, and cannot scale to the volumes required for competitive AI accelerator production. The gap at the frontier is not merely one of capability but of the entire manufacturing knowledge base that makes frontier production economically viable.

1.2 Capability Assessment: Maturity and Adoption Curves

The chokepoint architecture targets three distinct layers of the semiconductor supply chain. Understanding their relative maturity is essential to assessing where the controls bite hardest.

Supply Chain Layer	Key Controlled Items	China Domestic Alt.	Assessed Gap
EUV Lithography	ASML EUV scanners (NXE series)	SMEE — DUV only	**10–15 years**
Advanced Logic Chips	H100/H800 GPU class; sub-7nm	Cambricon, Biren — limited	**4–7 years**
EDA Software	Synopsys, Cadence toolchains	Huada, Empyrean — partial	**5–8 years**
Advanced Packaging	CoWoS, HBM integration	Progressing rapidly	**2–4 years**
Specialty Chemicals	High-purity photoresists	Partial substitution	**3–6 years**

The EUV lithography gap is the most durable chokepoint. ASML holds a monopoly on EUV scanner production, with no viable alternative technology in existence. China's domestic SMEE produces deep ultraviolet equipment that can, with multiple exposure passes, approximate 7nm production — but cannot approach 3nm. The physics of EUV — using plasma-generated 13.5nm wavelength light in precision vacuum chambers — represents accumulated engineering knowledge that took ASML over two decades to commercialize. China's path to EUV independence, even with unlimited resources, is measured in a decade or more.

1.3 Dependency Mapping: Where the Chokepoints Are

The semiconductor supply chain has four critical chokepoints where Western control is near-total and Chinese substitution is most constrained:

- **Chokepoint 1.** EUV lithography equipment — ASML monopoly; Dutch export license now required for all EUV shipments.
- **Chokepoint 2.** Advanced chip design software (EDA) — Synopsys and Cadence control ~70% of global advanced node EDA; both subject to BIS restrictions.
- **Chokepoint 3.** High-bandwidth memory (HBM) — required for AI accelerators; produced by SK Hynix and Samsung under U.S. export controls; China's domestic HBM capacity is negligible at required density.
- **Chokepoint 4.** Advanced packaging interconnect — CoWoS and 2.5D packaging critical for AI chip performance; TSMC controls most advanced capacity.

The dependency map reveals an important structural feature: the controls are mutually reinforcing. A Chinese firm that circumvents GPU controls and acquires H100-class chips still cannot domestically produce replacement chips when those chips fail or become obsolete. The controls do not need to be perfectly enforced in the short run because the long-run reproduction barrier — EUV — remains intact.

1.4 Risk Quantification: Disruption Probability Matrix

> *The controls are not a wall. They are a time tax — one that imposes compounding costs on Chinese AI capability development relative to a counterfactual in which the controls did not exist.*

The primary technology risk to the control architecture is not Chinese technical breakthrough but enforcement failure. Grey-market diversion through third-country intermediaries — particularly in Singapore, Malaysia, and the UAE — has allowed some restricted chips to reach Chinese end-users. The BIS added enforcement resources in 2023 and 2024, and expanded the Foreign Direct Product Rule to cover more downstream applications, but enforcement against sophisticated evasion networks remains imperfect. The risk is not that enforcement failure collapses the architecture, but that it reduces the time tax from a decade to perhaps seven or eight years.

Pillar II — National Security

2.1 Threat Landscape: The AI-Military Capability Feedback Loop

The national security logic animating the export control architecture is specific: advanced AI systems require advanced compute, advanced compute requires frontier semiconductors, and frontier semiconductors are the one input that Western policy can meaningfully constrain. The causal chain from chip access to military capability is already observable in China's AI-enabled defense modernization.

The People's Liberation Army has invested significantly in AI applications across ISR, autonomous systems, and command-and-control. These applications share a computational substrate: high-performance AI accelerators of precisely the class that BIS controls are designed to restrict. The honest assessment is that the controls slow but do not stop the relevant military applications. China entered the restriction period with a significant installed base of A100 and equivalent chips acquired before October 2022. The controls' military significance is primarily forward-looking: they constrain China's ability to scale to the next generation of AI-enabled capabilities — the ones that matter for the 2027–2035 strategic competition window.

2.2 Doctrine Analysis: Policy, Posture, and Intent

The October 2022 rules marked a doctrinal inflection point. The Biden administration explicitly abandoned the previous framework of allowing commercial technology exports as long as they did not have clear near-term military applications. The new framework — sometimes called the 'small yard, high fence' doctrine, though the yard has expanded considerably — holds that certain technology categories are strategically sensitive regardless of their immediate military relevance, because their long-term dual-use potential is too high to permit Chinese access.

The doctrinal tension that has not been resolved is between supply-side restriction and demand-side investment. The CHIPS and Science Act commits $52 billion to domestic semiconductor manufacturing and R&D — significant, but broadly assessed as insufficient to restore U.S. fab capacity to a leading position within the relevant strategic timeframe. The controls slow China;

the investments are meant to extend the U.S. lead. Whether the investment is scaled to the task is an open question.

2.3 Dual-Use Dimension: Where Commercial and Military Intersect

Technology	Commercial Application	Military Application
H100 / H800 GPU	LLM training, cloud inference	Autonomous targeting, signals intelligence
Advanced networking (400G+)	Hyperscale data center interconnect	C2 resilience, distributed sensor fusion
HBM3 / HBM3e memory	AI accelerator performance scaling	Real-time battlefield data processing
Advanced EDA software	Commercial chip design	Custom ASICs for military applications

The dual-use classification challenge has direct consequences for enforcement. Because the same H100-class chip that trains frontier AI models also enables ISR processing that matters for military applications, Chinese firms have strong commercial incentives to acquire these chips through intermediaries. The BIS enforcement challenge is not primarily about catching state actors; it is about disrupting the commercial supply chain that connects legitimate export with illicit end-use.

2.4 Escalation Modeling: Red Lines and Strategic Risk

The export control architecture introduces escalation dynamics that deserve explicit attention. China has characterized the controls as economic coercion and responded with restrictions on gallium, germanium, and graphite exports — materials used in semiconductor and defense manufacturing where China holds dominant market positions. This tit-for-tat dynamic has not escalated to a full trade war, but the trajectory is concerning.

The more significant escalation risk is military. U.S. defense planning around Taiwan operates on a 2027 timeline derived from PLA modernization assessments. If the export controls are understood in Beijing as an attempt to extend the window in which the U.S. can credibly defend Taiwan — which they partially are — then the controls themselves become a factor in Beijing's calculus about whether and when to act. This is the most consequential second-order effect of the architecture, and it is not adequately accounted for in public policy discourse.

Pillar III — Geopolitics

3.1 Power Distribution: The Multipolar Shift in Semiconductor Terms

The semiconductor export control regime is a case study in how a multipolar world produces asymmetric dependencies. The United States does not dominate semiconductor manufacturing — Taiwan and South Korea hold that position. What the U.S. holds is dominance in semiconductor design (Qualcomm, Nvidia, Intel, AMD), EDA software (Synopsys, Cadence), and manufacturing equipment (Applied Materials, Lam Research, KLA), alongside effective veto power over allied export policy. This is a different form of structural power than manufacturing dominance, but it is, for the specific purpose of restricting Chinese access to frontier chips, equally effective.

3.2 Alliance Architecture: The Trilateral Framework and Its Fault Lines

The chokepoint architecture requires trilateral coordination among the United States, the Netherlands, and Japan to function. These three countries collectively control the three most critical inputs that China cannot substitute domestically: U.S. chip design IP and EDA tools; Dutch EUV lithography equipment; and Japanese semiconductor-grade chemicals, legacy equipment, and specialty materials. The trilateral framework is functioning.

The fault lines are real but currently manageable. The Netherlands and Japan imposed restrictions under U.S. pressure, and both have domestic constituencies — ASML shareholders, Japanese equipment firms — that bear direct costs from forgone China sales. ASML has estimated that China represented approximately 46% of its revenue in 2023, a concentration that creates structural incentives to find accommodation as controls tighten.

South Korea presents a more significant fault line. Samsung and SK Hynix have massive memory manufacturing operations in China — SK Hynix's NAND facility in Wuxi, Samsung's Xi'an facility — which were granted validated end-user exemptions that have since been renewed annually. The effective exemption of Korean memory fabs in China from the full force of the controls reflects a real tension: full application would devastate Korean semiconductor firms' China operations without equivalent disruption to the Chinese advanced logic supply chain that is the primary target.

3.3 Economic Leverage: The Counter-Coercion Toolkit

China has responded with a counter-coercion toolkit targeting Western dependencies in critical materials. The 2023 restrictions on gallium and germanium exports — China controls roughly 80% and 60% of global production respectively — were a direct response to chip controls. The 2024 expansion to graphite continued this pattern. These restrictions have not yet caused supply disruptions of the severity China may have intended, because Western firms held inventory and accelerated alternative sourcing, but they have established a precedent for using Chinese material dominance as a retaliatory instrument.

3.4 Scenario Construction: Three Divergent Futures

The QAF requires explicit scenario construction rather than single-point forecasts. Three scenarios are presented for the 2025–2030 period, differentiated by the two variables with the highest leverage on outcomes: the durability of the trilateral control coalition, and the pace of Chinese domestic semiconductor progress.

Scenario A — Sustained Containment (Probability: 35%)

The trilateral framework holds through 2030 without significant defection. Enforcement improves incrementally, reducing but not eliminating grey-market diversion. China's domestic progress continues but remains constrained at the 5nm frontier by the EUV chokepoint. The military capability gap in AI-enabled systems is preserved through the 2027 window.

Scenario B — Partial Erosion (Probability: 50%)

One or more components of the trilateral framework weakens — most likely through a Dutch or Japanese government change, or a negotiated carve-out for memory manufacturing that expands over time. China achieves 5nm production using multi-exposure DUV techniques at limited yield by 2027–2028. The capability gap narrows but does not close. This is the most probable outcome.

Scenario C — Collapse and Acceleration (Probability: 15%)

The control framework fractures under geopolitical pressure — a Taiwan crisis forcing allied governments to recalculate costs — and a Chinese technical breakthrough in alternative lithography. China achieves sub-5nm production at commercial yield by 2029–2030. The military capability gap closes in the relevant AI-enabled domains.

QAF Intersection Analysis: Where the Pillars Converge

The cross-pillar synthesis is where the Quantingent Analytical Framework generates insight that single-pillar analysis cannot. Three intersections are of primary significance.

Intersection 1 — Technology × National Security: The Compute Threshold Problem

The technology analysis establishes that the EUV gap imposes a durable constraint on China's frontier chip production. The national security analysis establishes that the relevant military applications require frontier compute at scale. The intersection reveals a threshold problem policy has not adequately resolved: at what level of compute access does China's AI-enabled military capability become militarily decisive?

Military AI applications do not require frontier chips to be effective — they require sufficient chips at scale. China has sufficient installed base for current applications; the controls constrain next-generation scaling. The threshold question is whether the controls are calibrated to the right capability level, or whether they are addressing yesterday's concern while China's military AI programs find workarounds through architectural efficiency and larger installed bases of legacy hardware.

Intersection 2 — National Security × Geopolitics: The Alliance Cohesion Paradox

The national security case for the controls is strongest in U.S. terms. The geopolitical analysis reveals that the costs are borne disproportionately by U.S. allies: Dutch ASML, Japanese equipment firms, Korean memory manufacturers, and Taiwanese TSMC.

> *The U.S. benefits most from the controls. Its allies bear the most cost. This asymmetry is the primary structural threat to the architecture's durability.*

This asymmetry is the most underappreciated structural risk in the control architecture. The U.S. CHIPS Act subsidizes domestic semiconductor manufacturing, benefiting U.S. firms and workers. There is no equivalent allied industrial policy framework. If the controls are to be durable beyond the current political moment, some form of benefit-sharing mechanism is required to maintain allied commitment.

Intersection 3 — Technology × Geopolitics: The Self-Sufficiency Acceleration

The technology analysis establishes that Chinese domestic semiconductor progress is real but constrained. The geopolitical analysis establishes that China has identified semiconductor self-sufficiency as a top-tier strategic priority. The intersection reveals an acceleration dynamic: the controls are doing precisely what they are designed to do in the short run, and in doing so are creating the conditions that could make them less effective in the long run.

China's semiconductor investment has increased substantially in response to the controls. The controls have removed the commercial argument against heavy domestic investment in alternatives: Chinese firms no longer face the question of whether to invest in expensive domestic alternatives when cheap U.S. alternatives are available. The alternatives are no longer available. This means Western controls are now directly funding the acceleration of the Chinese semiconductor industry they are designed to constrain.

What to Watch: Named Indicators Through 2026

The QAF requires named, trackable indicators rather than general trend monitoring. The following will determine which scenario trajectory the chokepoint architecture follows:

- **Indicator 1.** ASML order book composition — specifically the China revenue percentage in quarterly filings. Any sustained increase above current levels indicates Dutch enforcement erosion.
- **Indicator 2.** SMIC N+2 process yield rates — any credible industry reporting on whether SMIC's advanced process yields are improving toward commercial viability.
- **Indicator 3.** BIS enforcement actions — frequency and target-type of export control enforcement actions is a direct measure of grey-market diversion scale.
- **Indicator 4.** Nvidia China-specific product SKUs — further product re-engineering by Nvidia signals continued demand pressure and regulatory arbitrage.
- **Indicator 5.** Korean fab exemption renewal — the annual renewal decision for Samsung and SK Hynix China operations is a leading indicator of U.S. willingness to sustain the full control architecture.

Conclusion

The chokepoint architecture is the most serious attempt to apply industrial policy as a strategic instrument in the post-Cold War era. It is working, in the sense that China's access to frontier semiconductors is materially constrained and its domestic alternatives remain years behind the frontier. It is also creating the conditions for its own eventual obsolescence, by driving Chinese investment in alternatives at a scale and urgency that would not exist without the controls.

The most important analytical conclusion is not about technology or national security in isolation, but about the intersection of alliance durability and time horizons. The controls need to hold for approximately a decade to generate the military balance outcomes that U.S. strategic planning depends on. The commercial and political pressures on the trilateral coalition are likely to intensify over that same period. Managing that tension — sustaining allied commitment while the costs of doing so remain high and the benefits accrue primarily to the United States — is the central strategic challenge of the chokepoint architecture.

QUANTINGENT will continue to track this topic through the Brief format, updating the named indicators identified above as data becomes available. The next Foundational Report will examine the defense implications of China's advances in quantum sensing and their effect on U.S. naval operational security.

Methodology Note

This report applies the Quantingent Analytical Framework (QAF v1.0) across all three pillars: Technology, National Security, and Geopolitics. Each pillar is analyzed through four structured layers before cross-pillar synthesis. All conclusions are drawn from open-source primary and secondary sources. Risk estimates represent analyst judgment, not quantitative modeling, and are stated with explicit uncertainty. Named indicators are selected for their discriminatory power — their ability to differentiate between divergent scenario trajectories — rather than their prominence in current news coverage.