

MLCCs — The Next HBM? THEMATIC

Rubin & Rubin Ultra to drive a step-function in per-rack MLCC content — ranking the global producers

A quiet 5¢ ceramic component is becoming one of the most supply-constrained, ASP-positive parts of the AI build-out — and tier-1 producers are pricing it accordingly.

What is an MLCC? — The 30-second primer

A **Multi-Layer Ceramic Capacitor (MLCC)** is a tiny passive electronic component — typically smaller than a grain of rice — that stores and releases electrical charge. Mechanically, it is a stack of alternating layers of ceramic dielectric material (usually barium titanate) and metal electrodes, fused together at high temperature. The more layers (up to 1,000+) and the thinner each layer (sub-1 micron in advanced parts), the higher the capacitance you can pack into a given footprint.

MLCCs do two things on a circuit board: **(i) smooth out voltage** by absorbing and releasing charge faster than the power supply can respond, and **(ii) filter noise** by shunting unwanted high-frequency signals to ground. On any given motherboard, GPU package, or server power module, you will find hundreds to thousands of them clustered around the chips they support — the more current the chip draws and the faster it switches, the more MLCCs you need around it. They are individually cheap (fractions of a cent each), collectively ubiquitous (one trillion-plus produced globally per year), and increasingly the chokepoint in AI server builds.

The Top-Down Thesis

A traditionally commoditised passive component is being repriced as a strategic AI input — driven by content explosion, capacity allocation discipline, and structural ASP support from the tier-1 oligopoly.

Content per rack is exploding, not just per server. An AI server uses 10–15× the MLCCs of a general-purpose server, and a single Nvidia GB300 NVL72 rack now requires ~440,000 MLCCs (Morgan Stanley channel checks). Vera Rubin VR200 NVL72 is estimated at ~600,000 (+36%), with Rubin Ultra (Kyber, NVL576) expected to push this materially higher as GPU TDP rises to ~1,800W per unit and switching frequencies climb.

Architectural substitution is amplifying the volume signal. Higher power demands are forcing a wholesale shift from aluminium polymer capacitors to MLCCs at the VRM/decoupling stages — SEMCO has explicitly flagged this as the dominant architectural trend in AI server power. This is a one-way trade-up: lower ESR, smaller footprint, and most importantly a multi-vendor supply chain (versus single-vendor polymer caps).

The spec is migrating up the value curve. Rubin-class power delivery requires *high-capacitance, low-ESL, ultra-small form factor* parts — Murata's 47µF 0402 MLCC (world-first, July 2025) is the canonical example. These are 5–10× ASP of commodity 0402 parts and the manufacturing yield is materially lower, both of which compound the value capture for tier-1 producers.

Supply is structurally tight and the tier-1s are pricing it. Lead times on high-cap / low-ESL grades have stretched from 8–12 weeks (late 2024) to 26–40 weeks (2026). Murata implemented its first price hike in three years (+15–35% on AI-server and auto grades, eff. April 2026); Taiyo Yuden raised +6–13% (eff. May); Samsung Electro-Mechanics is targeting double-digit hikes; Yageo and Walsin have followed in Taiwan. SEMCO's 2026 capex is set to more than double YoY, explicitly for AI MLCC and FCBGA.

Oligopoly economics protect the trade. The top 5 (Murata, SEMCO, Taiyo Yuden, TDK, Yageo) control ~68% of the high-end MLCC market; Murata alone holds an estimated ~70% of AI server-grade MLCC supply per industry channel checks. Customer qualification cycles run 12–18 months, erecting hard barriers to share migration even as Chinese capacity scales in commodity grades.

This is the HBM analogue in passives. The setup mirrors HBM in 2023 — a quietly bottlenecked, design-in-protected component sitting one layer below the headline AI chip, where content growth and ASP compound through the cycle. The window to position ahead of consensus is finite; Rubin VR200 ships H2 2026.

A note on valuation: *Several names in this report have begun re-rating, and we acknowledge the initial move. Our conviction, however, is forward-looking: it is the combination of share price momentum validating the thesis and earnings compounding as ASP hikes and volume mix flow through to margins that we expect to propel further gains. The re-rating is a confirmation that the market is beginning to see what we see and that the earnings inflection is still ahead.*

Why the Market Is Still Mispricing This

BURIED IN INDUSTRIALS COVERAGE	CYCLE PATTERN RECOGNITION	INDEX / FX OPTICAL DRAG
MLCC names are covered by industrials/electronics analysts, not semis or AI infrastructure desks. The AI content-growth narrative has not transmitted into consensus estimates, particularly on FY27 ASP assumptions.	Investors anchored to the 2018–2019 and 2022–2023 MLCC down-cycles (consumer / smartphone-driven) are underweighting the secular shift toward AI / auto, which now represents the highest-margin slice of the mix.	Japanese names (Murata, Taiyo Yuden, TDK) carry JPY translation noise and modest index sponsorship; Korean and Taiwanese names sit outside many global tech benchmarks — creating a true alpha pocket vs. consensus AI exposure.

The Content Math — MLCCs Per Rack Across The Nvidia Roadmap

PLATFORM	GENERATION	GPU TDP (W)	GPUS / RACK	MLCCS / RACK	VS PRIOR GEN	VS GB200	SHIP WINDOW
HGX H100 / H200 Server	Hopper	700	8 / server	~3,000–4,000 / server	—	baseline (server)	2023–24
GB200 NVL72	Blackwell	1,200	72	~30,000–40,000	~10x (server→rack)	1.0x	2024–25
GB300 NVL72	Blackwell Ultra	1,400	72	~440,000	~11–14x	~11–14x	2025–26
Vera Rubin VR200 NVL72	Rubin	~1,800	72	~600,000	+36%	~15–20x	H2 2026
Vera Rubin Ultra NVL576 (Kyber)	Rubin Ultra	~1,800+	576	~3.0–3.5m (LC est.)	~5x	~80–100x	2027

Sources: Morgan Stanley channel notes, TrendForce, Findchips supply tracker, SEMCO IR, Nvidia GTC 2025 disclosures, LC IC estimates. GPU counts per Nvidia OCP specifications. Rubin Ultra MLCC estimate is LC base case extrapolation from per-GPU MLCC density implied by GB300/Rubin disclosures applied to 8x GPU count — to be validated against supply chain disclosures closer to launch.

Why The Content Is Going Up — Five Independent Drivers

- (1) GPU TDP rising faster than voltage can drop. Rubin GPUs run at ~1,800W on sub-1V rails, meaning >1,800A of current at the package — every 100A of incremental current draw requires meaningfully more decoupling capacitance close to the chip, scaling roughly linearly with TDP.
- (2) Aluminium polymer cap → MLCC substitution. SEMCO and Murata have publicly flagged that high-CV mid-to-large MLCCs are replacing polymer caps at the VRM input stage. This is a 1.5–2x volume uplift on a per-board basis, independent of GPU performance scaling.
- (3) NVLink switch trays and CX-9 NICs. Each 9-switch NVLink tray plus ConnectX-9 800G/1.6T networking introduces its own dense MLCC clusters for SerDes power filtering — a meaningful incremental driver beyond the GPU compute trays themselves.
- (4) Embedded / landside MLCCs inside the package. Next-gen Rubin packaging integrates MLCCs *inside* the substrate (landside / embedded) to reduce loop inductance — an entirely new SKU category at premium ASPs which doesn't displace board-level MLCC count.
- (5) Higher rack density / NVL144 / NVL576. Nvidia's roadmap (NVL72 → NVL144 → NVL576 Kyber) multiplies GPU count per rack — Rubin Ultra Kyber alone is 8x the GPU density of NVL72, with corresponding multiplication of decoupling requirements.

Supply-Side: The Squeeze Is Already Showing

SIGNAL	LATE 2024	Q1 2026	IMPLICATION
Lead times — high-cap, low-ESL grades	8–12 weeks	26–40 weeks	Severe allocation; supports ASP
Murata price action	Flat (3 years)	+15–35% Apr 2026	First hike since 2023
Taiyo Yuden price action	Declining	+6–13% May 2026	Cycle inflection confirmed
SEMCO 2026 capex (vs 2025)	KRW 1.15tn	>KRW 2.0tn	>2x — explicit AI MLCC build
Murata AI-server MLCC share (est.)	~60%	~70%	Tier-1 design-in lock-in
B-B ratio (top Japanese names)	~0.9	>1.0	Demand > supply

Scoring Framework — Three Pillars, Each Scored 1–5

The thematic case is broad, but capturing it through the right names requires discrimination. A pure-play, technologically advantaged, well-qualified producer will see a multi-x earnings impact from the content explosion; a diversified conglomerate with a small MLCC business will see only a marginal lift even if its absolute MLCC dollars rise meaningfully. We score each player across three orthogonal pillars and take a simple average for the composite ranking. Pillar selection and rationale below:

PILLAR	WHAT WE MEASURE	WHY IT MATTERS
Technology / IP (1–5)	Ability to manufacture ultra-small (0201/0402), ultra-high-capacitance ($\geq 10\mu\text{F}$, up to $47\mu\text{F}$) MLCCs with low-ESL/ESR characteristics required for sub-1V GPU decoupling. Bonus credit for in-house dielectric powder, embedded / landside MLCC capability, and proprietary thin-film dielectric process.	The Rubin-class AI server spec is at the technological frontier of MLCC manufacturing. Only 3–4 producers globally can deliver the leading-edge specs at production yields; this is the highest-margin slice of the mix and the hardest to displace. Pure proxy for share of the high-end (versus commodity) MLCC market.
Qualification (1–5)	Design-in / qualified status at Nvidia for GB200, GB300, and Rubin platforms; presence in BOMs of tier-1 ODMs (Foxconn, Quanta, Wistron, Inventec, Supermicro); existing automotive AEC-Q200 qualifications as proxy for high-reliability process capability.	Qualification cycles run 12–18 months and customer audit costs are prohibitive. A name not already qualified on Rubin platforms cannot meaningfully participate in the H2 2026 ramp regardless of technical capability. This is the gating factor for converting capacity into revenue.
Purity (1–5)	MLCC revenue as % of group revenue. Higher concentration = higher earnings beta to the thematic. Scored on the following rough bands: $>60\% = 5$; $40\text{--}60\% = 4$; $25\text{--}40\% = 3$; $10\text{--}25\% = 2$; $<10\% = 1$.	The thesis is sector-specific; we want the share-price impact undiluted by unrelated business lines. A diversified industrial conglomerate may have a top-tier MLCC technology but only see 5–10% EPS uplift; a pure-play sees 30–50%+ earnings beta to the same revenue impulse. Purity is the leverage multiplier on the trade.

Why These Three And Not Others — Pillars We Considered And Excluded

Capacity / scale. Considered but excluded — capacity is a function of the three pillars above (technology gates which capacity is relevant; qualification gates which capacity is monetisable; purity captures the share of capacity going to the thematic). Adding it would double-count.

Geographic risk / sovereign exposure. Considered but excluded from the composite — handled as a separate portfolio-construction overlay rather than a scoring axis. Taiwan-domicile concentration and China-customer exposure are real risks but should be sized at the portfolio level, not used to handicap individual security selection.

Valuation. Deliberately excluded from the ranking. The ranking is a quality / exposure screen; valuation is the second step in positioning and is handled in the deep-dive notes to follow. Building valuation into the screen would conflate "what to own" with "what to pay" — they answer different questions.

ESG / governance. Acknowledged as a relevant overlay (particularly for Chinese names) but does not drive thematic selection at this layer.

Composite & Decision Rules

Composite score = simple average of the three pillars, rounded to one decimal. We deliberately use an equal weight rather than a weighted average: each pillar is a distinct gating condition, not a proportional contributor. A score of 5 in technology with a 2 in qualification is essentially un-investable for this thesis regardless of average; the simple-average frame surfaces these imbalances clearly in the underlying scores.

AI-server torque is reported as a qualitative tag alongside the composite — it indicates the directional sensitivity of group earnings to incremental AI-server MLCC demand, accounting for both the purity and the qualification scores. This is the tag we use to size positions, not the composite alone.

Ranking tie-breaks are resolved in favour of higher purity, on the rationale that pure-play exposure is the dominant leverage variable in any thematic basket. Where purity is tied, qualification breaks the tie.

The Top 10 — Global MLCC Producers Ranked For The Rubin Content Thesis

RANK	COMPANY	TICKER	DOMICILE	MLCC REV SHARE	TECH / IP	QUAL	PURITY	COMPOSITE	AI TORQUE
1	Taiyo Yuden	6976 JP	Japan	~64%	4	5	5	4.7	Very High
2	Samsung Electro-Mechanics	009150 KS	South Korea	~45%	5	5	4	4.7	Very High
3	Murata Manufacturing	6981 JP	Japan	~40–45%	5	5	4	4.7	High
4	Yageo (incl. KEMET)	2327 TT	Taiwan	~40%	4	4	4	4.0	High
5	Walsin Technology	2492 TT	Taiwan	~50%	4	4	4	4.0	Medium-High
6	Prosperity Dielectrics (PDC)	6533 TT	Taiwan	~65%*	4	3	5	4.0	High (niche)
7	TDK Corporation	6762 JP	Japan	~10–12%	5	4	2	3.7	Low (diluted)
8	Kyocera (AVX)	6971 JP	Japan	~8–10%	4	4	2	3.3	Low (diluted)
9	Samwha Capacitor	001820 KS	South Korea	~55%	3	3	4	3.3	Medium
10	Holy Stone Enterprise	3026 TT	Taiwan	~55%	3	3	4	3.3	Medium

* PDC: capacitors + dielectric powder combined; powder is sold to other MLCC producers as well. MLCC revenue shares are LC estimates from latest available company filings and industry reports (TrendForce, Mordor Intelligence, BigGo, Digitimes channel checks). Highlighted rows = top-3 conviction names. Composite = simple average of three pillars rounded to one decimal. Maruwa (5344 JP) intentionally excluded — primarily a ceramic substrate (AlN/Si3N4) play, MLCC is a small ancillary line; should be evaluated under power-semis ceramic substrates theme separately.

Top-3 Per-Name Conviction Snapshot

NAME	WHY IT RANKS HERE
Taiyo Yuden — #1	The purest-leverage play in the universe. MLCC = ~64% of group revenue, the highest of any tier-1. Smaller absolute MLCC market share (~10–13%) but B-B ratio >1, capacity running flat-out, and first to push price (+6–13% eff. May 2026). Highest earnings beta to MLCC ASPs in the entire list — every 5% ASP move translates near-linearly into EPS. The "if the thematic works, this is the maximum slope" name.
SEMCO — #2	Highest-conviction primary pick. Q1 2026 call: 2026 capex more than doubling YoY, explicitly directed at AI-server high-cap MLCC + FCBGA + silicon capacitors. Just announced KRW 1.6tn AI chip component contract with US big tech (May 2026). Third Philippines plant under construction. SEMCO is leading the industry narrative on aluminium polymer-to-MLCC substitution and has the embedded MLCC roadmap — the cleanest combination of technology, design-win momentum, and pure-play exposure.
Murata — #3	The quality anchor. ~40% global MLCC share and an estimated ~70% of AI server-grade supply. World-first 47µF in 0402 (July 2025) — the canonical Rubin-class spec. First price hike in 3 years at +15–35% (eff. April 2026). Ranks #3 only on purity (~40% MLCC, the rest is inductors / filters / EMI — all of which also benefit from AI but dilute the per-share leverage). Lowest beta, highest quality; the position you can size meaningfully without single-stock risk.

Suggested Portfolio Construction

Core (~60% of basket): Murata (#3) and SEMCO (#2) as the quality anchors — large-cap, liquid, dominant share, and price-setting power. **Beta sleeve (~25%):** Taiyo Yuden (#1) for the pure-play ASP leverage. **Taiwan diversification (~10%):** Yageo (#4) or Walsin (#5) for order-transfer benefits as Japanese/Korean tier-1s ration capacity to marquee customers. **Small-cap optionality (~5%):** PDC (#6) on the dielectric-powder chokepoint angle — illiquid, watchlist-sized only. The TDK and Kyocera dilution penalty makes them inferior expressions of this specific thesis despite their technical capability; better held for other reasons.

The above sizing is illustrative for a single-theme basket and assumes an active equities sleeve. Sizing within an existing portfolio mandate must be calibrated against the client's existing AI infrastructure exposure (Nvidia, hyperscalers, semiconductor equipment, power & cooling names) to avoid stacking correlated factor risk. **This is not investment advice or a client suitability assessment; figures are based on publicly available industry data and channel checks, some of which are estimates and have not been independently verified against Bloomberg fundamentals. Any client distribution requires Investments & Compliance review.**

LIGHTHOUSE CANTON

Singapore

16 Collyer Quay #11-02
Collyer Quay Centre
Singapore 049318

☎ +65 67130570

India

Unit No-104A,Worldmark
2 Asset Delhi Aerocity,
New Delhi 110037

☎ +91 9650473961

Unit No 507/508, A Wing,
INS Tower, G Block, BKC,
Mumbai- 400051

☎ +91 9650473961

UAE

The Exchange Gate Village 11,
Unit 802 Dubai International
Financial Centre PO Box 507026
Dubai, UAE

☎ +971 45 861500

1st Floor, WeWork37,
Cunningham Cross Rd,
SRT Road Vasant Nagar,
Bengaluru-560001

☎ +91 9900096873

RK Swamy Centre, Hansa
Building, Door No:3, Thousand
Lights,
Chennai-600006

☎ +91 9650473961

UK

24 Hanover Square,
London W1S 1JD

☎ +44 164 2843 487

Suite 502, Building 450, Central
Plaza, Genome Valley,
Shameerpet,
Hyderabad 500 078

☎ +91 9650473961

Unit No FF-10, FF Floor, Pragya
Accelerator, Block 15T GIFT CITY,
Gandhinagar
Gujarat-382355

☎ +91 9650473961

✉ info@lighthouse-canton.com

📌 Lighthouse Canton

DISCLAIMER

The contents of this document are confidential and are meant for the intended recipient only. If you are not the intended recipient, please delete all copies of this document and notify the sender immediately.

This document, provided as a general commentary, is for informational purposes only and is not to be construed as an offer to sell or solicit an offer to buy any financial instruments in any jurisdiction. This does not constitute any form of regulated financial advice, and your independent financial advisor should be consulted prior to taking any investment decision(s).

This document is based on information from Sources which are reliable but has not been independently verified by Lighthouse Canton Pte Ltd and its affiliate companies ("LC"). LC has taken the reasonable steps to verify the contents of this document and accept no liability for any loss arising from the use of any information contained herein. Please also note that past performances are not indicative of future performance.

Information contained herein are those of the author(s) and does not represent the views held by other parties. LC is also under no obligation to update you on any changes made to this document.

This document is prepared by Lighthouse Canton Pte Ltd and its affiliate company, Lighthouse Canton Capital (DIFC) Pte Ltd, which are regulated by Monetary Authority of Singapore ("MAS") and Dubai Financial Services Authority ("DFSA") respectively. MAS and DFSA has no responsibility for reviewing, verifying and approving the contents of this document and/or other associated documents. The contents of this document may not be reproduced or referenced, either in part or in full, without prior written permission from LC.

This document is confidential and is only intended for Accredited Investors and/or Professional Clients, as defined by MAS and DFSA.