NORTHCAPE
CAPITAL

Global Emerging Markets

AI and the Changing Face of the EM Index

The EM index is no longer providing investors with diversified exposure - it's becoming an AI bet in disguise. This article unpacks the dangerous concentration, rising exuberance, and why active investors may find better opportunities beyond the market's hottest trade.

Ten years ago (in 2016), the three largest stocks in the EM index had a combined weight of less than 10% and the top ten index constituents had a combined weight of less than 20%. Seven countries were represented amongst the top fifteen index positions. Today, the top three index members have a combined weight of almost 25%, and the top ten stocks make up more than a third of the entire index (see Exhibit 1). Only four countries (Taiwan, South Korea, China and India) are represented amongst the fifteen largest companies, in an index which spans 24 countries and is supposed to represent a diverse set of economies that contribute around half of total global GDP. Consequently, the index no longer provides either a true reflection of the underlying drivers of emerging markets nor sadly adequate diversification for investors.

Exhibit 1: Increasing MSCI EM Index Concentration 2016 to 2026, 30 April Year End (Source: Bloomberg)

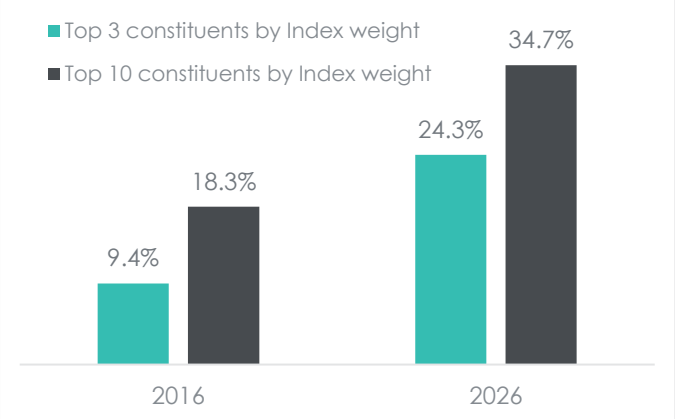

Related to the alarming increased concentration of the EM index is overexposure to a single investment theme: artificial intelligence (AI), and more specifically the frenzied global build out of AI data centres. The three largest index members (TSMC, Samsung Electronics, and SK Hynix), which now comprise almost one quarter of the entire EM index, are all semiconductor manufacturers. Rather than broad demand for these companies' products, recent results and share price trading has been overwhelmingly influenced by a single driver: the construction of AI data centres and associated booming demand for AI semiconductor chips. Currently the seven largest stocks in the index, which have a combined weight of almost a third of the whole benchmark, are all effectively investment plays on the same theme. More specifically, if we observe the weighting of the top ten stocks in the EM index, we note the top seven are all very tightly linked to the AI investment narrative with a weighting of 32.2%, up from the 10.3% weighting these ten stocks had in 2016 – a massive increase in concentration based exclusively on the AI phenomenon (see Exhibit 2).

Exhibit 2: MSCI EM Index Top 10 companies 2016 v 2026 at 30 April (Source: MSCI)

Current Top 10 - MSCI EM Index	Country	Rank	Sector	2016	2026
TSMC	Taiwan	1	IT, AI-Linked	3.1%	14.2%
Samsung Electronics	Sth Korea	2	IT, AI-Linked	3.3%	6.0%
SK hynix	Sth Korea	3	IT, AI-Linked	0.4%	4.1%
Tencent	China	4	IT, AI-Linked	2.9%	3.3%
Alibaba	China	5	IT, AI-Linked	0.0%	2.4%
Delta Electronics	Taiwan	6	IT, AI-Linked	0.2%	1.1%
MediaTek	Taiwan	7	IT, AI-Linked	0.3%	1.1%
China Construction Bank	China	8	Financial	1.5%	0.9%
HDFC Bank	India	9	Financial	0.0%	0.8%
Reliance Industries	India	10	Energy	0.5%	0.8%
Total Top 10 Weighting In EM Index				**12.3%**	**34.7%**
Total Top 10 AI-Linked Weighting In EM Index				**10.3%**	**32.2%**

In addition to the lack of diversification the index consequently now provides, there are worrying signs of irrational exuberance in the valuation of many AI related stocks. Delta Electronics, a top ten MSCI EM index member, trades on approximately 55x P/E whilst its listed Thai subsidiary trades on more than 100x P/E. Both are more than double the long-term average multiples. Similarly, SK Hynix, the third largest stock in the index, at close to 10x price/book currently trades on more than five times its long-term average price/book ratio.

Investors will note in November 2025 we reduced our exposure to this sector, considering the valuation risks, especially relative to position sizing, and growing concerns about some mean reversion in semiconductor pricing as supply eventually catches up with demand. However, what if we are wrong? What are the counter arguments to suggest that the AI investment theme has a lot further to run, and demand for AI semiconductors and related IT services will outstrip supply for much longer than we anticipate?

Fly on the wall: extracts from our Devil's Advocate session on AI

As part of our investment process, we continuously retest our current views in light of new information, and this is integrated through our appropriately named "Devil's Advocate" sessions. The magnitude of the semiconductor sector's outperformance has provided a useful backdrop for these debates, prompting us to revisit key elements of our bear thesis – namely that AI capital expenditure will normalise and that these stocks will ultimately return to earth.

Whilst we would be attempting false precision if we tried to calculate equilibrium demand, supply and pricing for AI, we could take note of structural differences between the impact of AI capex spending versus historical levels. The purpose of this is to separate companies that should be permanently rerated from companies that are over-earning and are benefitting from markets flush with hot money chasing the AI narrative.

Devil's advocate statement #1: AI capital expenditure permanently changes the demand limitations previously constraining the semiconductor cycle.

Before the AI boom, the semiconductor industry was constrained by structural forces. The industry ran on relatively short-term boom/bust cycles where capacity additions would overshoot demand, crushing ASPs and vice versa. Moore's Law obligated chipmakers to deliver more compute per dollar in each cycle and the few key customers commanded volume leverage.

Broadly speaking, there was a tangible demand ceiling. A finite number of PCs, cloud servers and mobile phones could be sold each year, and this broadly matched the tangible volume of chips that could be manufactured annually. These chips had a price ceiling as a cost of goods sold, or cost-plus pricing structure within a product that had a replacement lifecycle. To over-simplify a complicated sector, the semiconductor industry was cyclical but constrained by fundamental economic principles where supply and performance gains impacted pricing and end-market demand from consumers on the back of advancements in technology – in short, the semiconductor industry was a price taker.

AI has completely broken down this model. At present, AI demand has no tangible ceiling, no visible saturation point and each generation of model requires more compute to train and serve (hence the "exuberance" in the market). New generations of models introduce more capabilities, and industries

are being forced to adopt AI to remain competitive, creating the feedback loop of requiring even more compute to train and serve. Additionally, whilst we cannot forecast the demand ceilings, customers (largely the "hyperscalers" with strong balance sheets) have provided visibility into the hundreds of billions of dollars that they intend to spend on AI infrastructure for at least the next 24 months.

The change in the nature of the demand, and the surety of which this exuberance continues has led capital markets to extrapolate current conditions into the future, as opposed to anchoring back to mean-reversion. Importantly, this exuberant, irrational demand must plateau once enterprises are satisfied with the level of artificial intelligence that can be economically integrated into their operations in a viable, and value-additive manner. At this point demand growth for AI chips starts to slow.

However, in the short-term, the market's view on demand for semiconductors has clearly shifted from "finite" to "theoretically infinite" with AI applications. This theory, combined with tight supply of semiconductor chips, is seeing large amounts of capital flows rush into this asset class, supporting the rapid rise in AI linked stock valuations.

Devil's advocate statement #2: Supply scarcity is a more durable condition, not a temporary undersupply

With what is widely assumed to be insatiable, well-funded demand for AI infrastructure, we turn to the supply of its key components. As previously alluded to, the supply of components loosely matched the demand for components in pre-AI devices. The new demand not only outstrips the supply but does so in a different ratio. This is because AI workloads use different ratios of compute resources: GPUs, CPUs, memory (HBM) and components that connect these chips together.

Training is GPU-dominated, whereas inference is more balanced across GPU, CPU and HBM. Because there is large scale usage of AI across enterprises, heavy inference usage has caught out the memory sector as being grossly undersupplied.

These components are manufactured in the leading-edge foundries, or the most technically advanced foundries available. Expanding this capacity is bound by several chokepoints that are held by few players. For example, the bottlenecks are layered where there is a single provider of Extreme Ultraviolet (EUV) tools, few providers of advanced packaging capacity, and HBM DRAM comes from just three providers (SK Hynix, Samsung and Micron) at scale. All these individual components operate in areas so advanced, the degree of precision is almost impossible to imagine.

Now this is not structurally different from prior cycles where fab capacity additions eventually resolved shortages, but what is different is that we are seeing memory players (Samsung, SK Hynix, TSMC) gaining significantly more pricing power as they control very tight supply, and barriers to entry are higher - in fact they may have moved from being "price takers" to "price makers" in the short-term. As such, the durability of this supply scarcity has shifted more bargaining power to memory players. This has led to enormous increases in pricing for semiconductor chips, in some cases contract prices are up 5-10x over the past six months, which in turn has given rise to 400-500% lifts in operating profits for the leading semiconductor companies over the same time frame.

We have also seen evidence of changes in customer contracts, shifting towards longer-term agreements (LTAs). This has added to the argument that increased chip volumes with highly elevated pricing should support sustainably much higher earnings for a lot longer than the past. This in turn deserves a materially higher multiple, which is exactly what has happened to these stocks – a massive rerating up resulting in a much higher index weighting for the sector.

Devil's advocate statement #3: Some, but not all AI is genuinely useful and will generate ROIC above WACC on its AI capex at maturity.

A key datapoint over April was the impressive growth in revenue from OpenAI and Anthropic over the past three months, with the two models alone generating ~USD$60 billion in annual recurring revenue (ARR) over the month of March, with no signs of slowing down. In fact, revenues from these two are forecast to increase by nearly 4x to over US$230 billion by December 2026! (see Exhibit 3).

Additionally, throughout earnings season we have been seeing large enterprises adopting AI with demonstrable cost efficiencies. For example, Bank Central Asia of Indonesia outlined meaningful improvements in efficiency, customer experience, and fraud detection from AI which all structurally improve its cost-to-income ratio. Walmex (Mexican grocery retailer) unveiled AI-driven real-time, granular inventory tracking resulting in significant SKU optimisation. As these models improve and purpose-built agentic AI is adopted by enterprises, high growth monetisation scenarios remain believable.

Pricing models for AI are in nascent form but we are seeing early indicators of *value-based pricing (such as percentage of a salary saved via AI)*. Adoption in this space has been low but can be economically rational for parties when outcomes are verifiable and material. Whilst there is not yet a dominant form of pricing models in AI, we are seeing examples of pricing tied to a business outcome which makes the Total Addressable Market (TAM) potentially huge, and frustratingly more difficult to quantify.

With Agentic AI appearing to have crossed several inflection points, one could believe it is directionally heading into the area of being priced on value as opposed to cost-plus. However, if one believes that software and hardware improvements will continue to reduce the cost of generation, based on this premise there will be an eventual commoditisation of AI pricing, which means value-based pricing is unlikely to be sustained long-term.

Meanwhile in the short-term, the rapid pace of AI adoption has stretched demand with the super cycle feeding its own self-fulfilling prophecy. This long-drawn out chokepoint means that AI semiconductor stocks could be as hot as, well, NVIDIA in 2023-2025 after the use-case of GPUs changed from gaming to high-performance computing. Enter triple-bull leveraged Memory ETFs to retail traders for 2026!

Exhibit 3: GenAI Annual Recurring Revenue (US$bn) 2026 YTD and Estimated (Source: Arete Research)

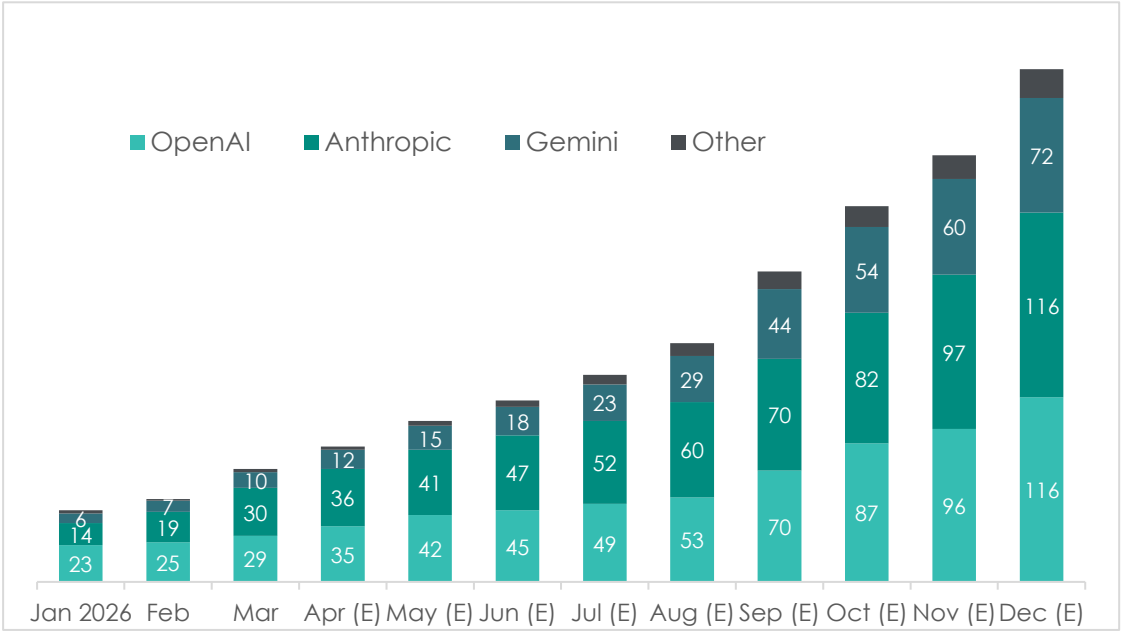

Signs of Exuberance

Our Devil's Advocate analysis highlights some of the fundamentals and theories driving the rapid growth in demand for AI related IT components (especially semiconductor chips). This does support a fundamental improvement in the valuation of these stocks, no question.

However, from a capital markets perspective we are seeing several signs that suggest the "AI theme" is attracting speculative flows and instruments (such as triple-bull leveraged Memory ETFs). Share prices are not being purely driven by the fundamental long-term outlook for the sector.

One example at a fundamental level, is the recent run up in Taiwanese Printed Circuit Board (PCB) manufacturers' share prices. PCB is historically a low margin commoditised product with many manufacturers competing on cost and price. It is of no coincidence that the industry has largely migrated to China, where large government subsidies have driven Chinese manufacturers' competitive advantage.

Exhibit 4: Major Taiwanese PCB manufacturers share price performance (Source: Bloomberg)

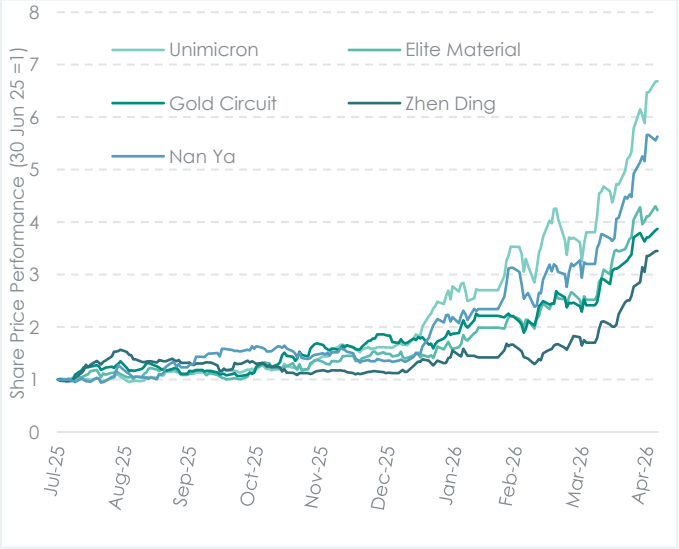

Since January 2026, several Taiwanese PCB manufacturers have seen their share price appreciate 2-3x in a matter of months (see Exhibit 4). The recent run up has been driven by news that NVIDIA will include a more advanced form of PCB in its AI server products. This has caused a significant increase in expectations for PCB industry profitability over the coming years. However, there also appears to be a significant increase in industry capacity with capex from just the major listed firms (many Chinese PCB manufacturers are not listed entities) doubling between 2024 and 2026 (see Exhibit 5).

Exhibit 5: Major Taiwanese and Chinese PCB Manufacturers' Capex (US$m) (Source: Bloomberg)

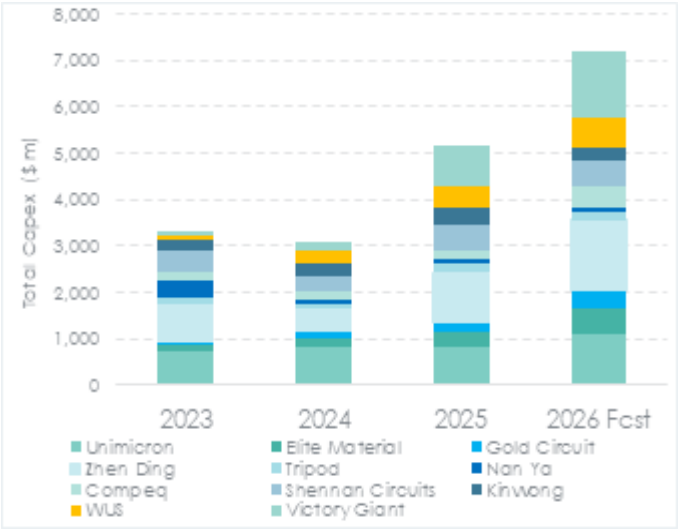

We note that PCB technology does not require large, sophisticated plants like semiconductors, so installing capacity has very limited lead times (it takes just a few months). And PCB company profitability is still driven by the laws of demand and supply. To us, the rise in PCB share prices and valuations (the price-to-book ratios of Taiwanese PCB manufacturers is currently averaging 10x) is not considering the economic reality of how this industry operates over the medium term.

We are also witnessing growing ETF inflows into the semiconductor sectors in Korea and Taiwan (see Exhibits 6 and 7). We note that these flows have increased materially in 2026 following very strong share price performances of key AI stocks in 2025. On a 12-month basis, the share price increases of these companies are now enormous: TSMC (+151%), Delta Electronics (+538%), SK Hynix (+744%), and Samsung (+387%).

Exhibit 6: Monthly Inflows in Semiconductor Focussed ETFs (US$m) Source: Bloomberg

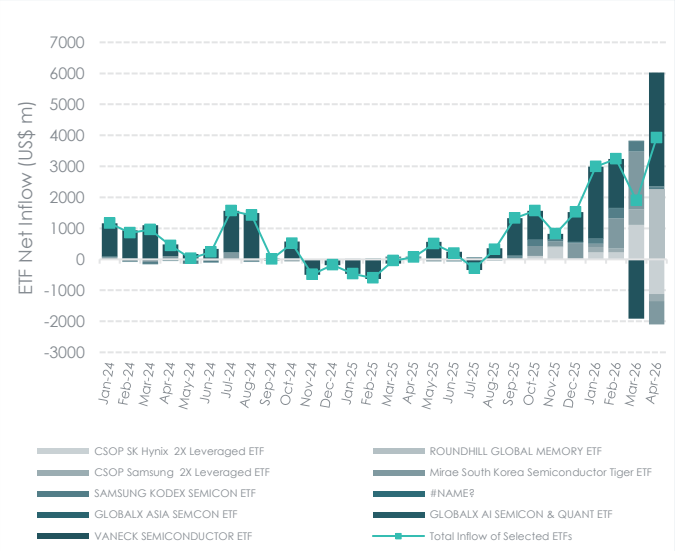

Exhibit 7: Monthly Inflows in Korean and Taiwanese Domiciled ETFs (US$m) Source: Bloomberg

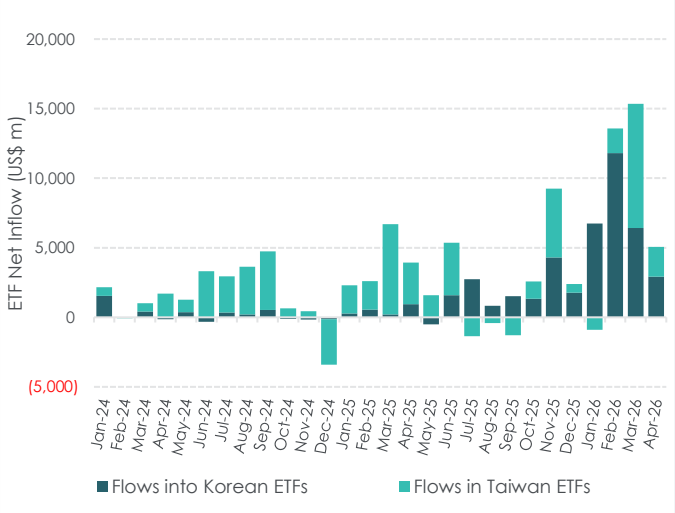

Summary Implications of AI boom to EM portfolio construction

The change in structure of the EM equity asset class over the past few years represents a real challenge to dedicated EM investors, especially those benchmarked to the index. As the foregoing analysis shows AI linked stocks in the top ten alone are now 32% of the index, and if we add in AI-linked stocks outside the top ten, this number will grow to be close to 40%.

To outperform the EM index over the past year, we would need to have had at least 40%, potentially as much as 50%, of our portfolio dedicated to the AI investment theme.

Frankly, from our perspective that is unacceptable risk to take. We prudentially reduced our exposure to the AI semiconductor sector in November 2025, as we saw bubble tendencies emerging and have maintained a weighting at about 20% of our total assets to this sector since then, well below Index. This weighting is still a reasonable absolute exposure reflecting many of the elements in our Devil's Advocate work, the counter view, which is tightly integrated into our investment process. The main point justifying our 20% portfolio ownership is the positive risk that demand for AI chips outstrips supply to such an extent over the next year or so, it creates a paradigm shift in semiconductor pricing. That is DRAM/NAND chip pricing migrates from short-term contracts/spot pricing (cyclical) to long term agreements (structural), thereby locking in these super profits for a lot longer than normal.

The other point to make (which we touched on in the EM Performance section) is the rapid rise in liquidity rushing into AI-linked stocks has in many cases been funded by indiscriminate selling of other stocks – some are on our EM Approval List. This is creating some exceptional bargains, which we have been adding to over the past six months. By way of an example, Bank Central Asia is now trading at its lowest multiple over the past twenty years, in line with its GFC nadir – and by the way it has huge potential to reduce costs and improve returns on capital by deploying AI – an opportunity the market is completely ignoring today!

Contrast BCA's valuation with SK Hynix – an AI semiconductor stock we have significantly reduced exposure to. SK Hynix is trading at a highly elevated P/B multiple of 9.4x – an extreme level of valuation, never seen in its entire history (see Exhibit 8).

Exhibit 8: The Yin and Yang – Bank Central Asia v SK Hynix (Price/Book) 2006-2016 (Source: Bloomberg)

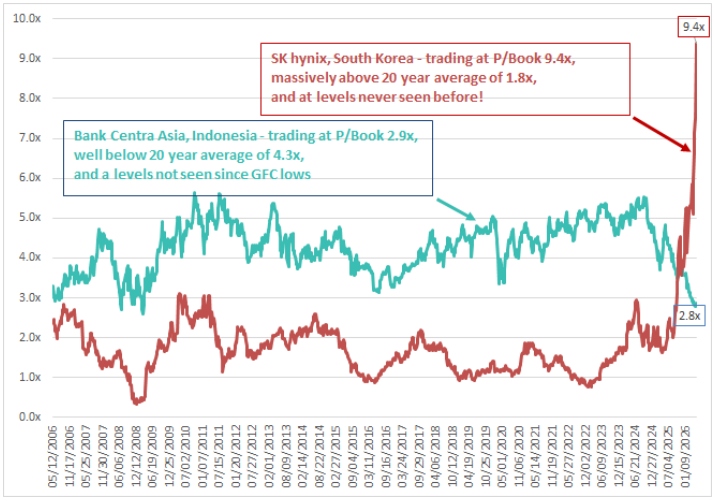

To reiterate, the EM index no longer provides either a true reflection of the underlying drivers of emerging markets, nor sadly adequate diversification for investors. This reinforces the need for active management in the EM equity asset class more than any other time in its history.

For more information, please contact our distribution team:

Katie Orsini
Client Director

p: +61 2 8234 3629
e: katie.orsini@northcape.com.au

Wendy Hammond
Client Manager

p: +61 2 2 8234 3638
e: wendy.hammond@northcape.com.au

Eliza Clarke
Client Manager

p: +61 2 2 8234 3606
e: eliza.clarke@northcape.com.au

Northcape Capital

Sydney
Level 24, 45 Clarence
Street Sydney NSW 2000
Australia
P: +61 2 8234 3600

Melbourne
Level 18, 90 Collins Street
Melbourne VIC 3000
Australia
P: +61 3 8626 8000

ABN: 53 106 390 247
AFSL: 281767

www.northcape.com.au