



Capital  
Markets

BANK

# The Software Investor's Handbook to Generative AI

RBC Imagine™

Thinking further forward

EQUITY RESEARCH | NOVEMBER 9, 2023

For Required Non-U.S. Analyst and Conflicts Disclosures, see page 16





Capital  
Markets

November 9, 2023

## RBC Imagine™: The Software Investor's Handbook to Generative AI

**Our view:** Following our flagship RBC Imagine™ report on GenAI (see [here](#)), we have continued to iterate on our framework for the implications of GenAI on software, most notably in our Software Download pieces (see the most recent one [here](#)).

In this note, we package all our thoughts and musings in one easy-to-read report (in many cases expanding on what we've previously written). Topics covered include:

- **Why we view GenAI as such a transformational technology**, and why we believe investors are overestimating the short-term impact, but underestimating the long-term impact of GenAI.
- **Our framework for evaluating GenAI names**, including a historical perspective, short-term vs. long-term implications, and opportunities where we think the market is “wrong”.
- **Four categories of GenAI beneficiaries**, including: 1) large incumbent vendors that can leverage their data and distribution to take advantage of the technology; 2) vertical software leaders that have not just significant data, but also the ability to verticalize GenAI solutions to accelerate the winner-take-all dynamics of vertical software; 3) mid-market challengers that can use GenAI as an innovation engine to narrow the gap with market leaders; and 4) companies that enables generative AI for others.
- **Four categories of companies at-risk from GenAI**, including: 1) legacy, on-premise software companies that cannot truly embrace GenAI, which are cloud-only services; 2) companies that don't fully embrace GenAI and rearchitect around it; 3) analytics companies claiming to be AI; and 4) companies whose base functionality could be potentially replaced by generative AI.
- **The different paths of GenAI monetization**, including direct (discrete monetization, paywalling, and consumption) and indirect (higher free-to-paid conversion, improving customer retention and unit economics), and the puts and takes of each path, as well as separating what is truly monetizable, versus what is tablestakes.
- **The margin implications of GenAI**, including the high costs of compute and our structural outlook for gross margins (we estimate GenAI workloads are currently 5x more expensive than traditional cloud workloads). Long term, we expect GenAI to be a ~60% gross margin business, but to be accretive to gross margins dollars (due to a long-term 2x-3x pricing uplift relative to cloud workloads) and neutral to FCF margins (with greater OpEx efficiencies from GenAI).
- **Our thoughts on the GenAI platform wars** (primarily OpenAI vs. Anthropic vs. Cohere), as well as digging into the debate on proprietary versus open source LLMs (especially LLaMa 2) and whether LLMs are commoditized.
- **The importance of Vector databases**, why they are necessary, and whether they are a feature or standalone platform.
- **Overcoming the hurdles to enterprise adoption of GenAI**, including costs, hallucinations, data privacy, and domain expertise.
- **Societal implications**, including impacts from governments and regulators, ethical, legal, and data privacy concerns.
- **Generative AI vs. Chat** and why ChatGPT is just a piece of the broader GenAI picture.
- **Putting this altogether, we hope this helps software investors** navigate the rapidly-evolving and increasingly complex GenAI landscape. We include a comprehensive index of relevant published research notes.

Our favorite names levered to GenAI include **ADBE, CFLT, COUR, CRWD, CVO, CWAN, ESTC, GTLB, GWRE, HUBS, MDB, MSFT, NET, NICE, NOW, PME, SHOP, WDAY, XRO, and ZM**, while the names we view most at risk from GenAI include **ASAN, FFIV, PLTR, SMAR, and ZI**.

RBC Capital Markets, LLC  
**Rishi Jaluria** (Analyst)  
 (415) 633-8798,  
[rishi.jaluria@rbccm.com](mailto:rishi.jaluria@rbccm.com)  
**Matthew Hedberg** (Analyst)  
 (612) 313-1293,  
[matthew.hedberg@rbccm.com](mailto:matthew.hedberg@rbccm.com)  
**Matthew Swanson** (Analyst)  
 (612) 313-1237,  
[matthew.swanson@rbccm.com](mailto:matthew.swanson@rbccm.com)  
**Richard Poland** (AVP)  
 (347) 920-6958,  
[richard.poland@rbccm.com](mailto:richard.poland@rbccm.com)

Royal Bank of Canada, Sydney Branch  
**Garry Sherriff** (Head of Australian Equity Research & Lead Technology Research)  
 +61 2 9033 3022,  
[garry.sherriff@rbccm.com](mailto:garry.sherriff@rbccm.com)

RBC Dominion Securities Inc.  
**Paul Treiber**, CFA (Analyst)  
 (416) 842-7811,  
[paul.treiber@rbccm.com](mailto:paul.treiber@rbccm.com)



**RBC Imagine™**  
 Thinking further forward



## Table of contents

Why GenAI is so important.....	4
Framework for evaluating GenAI benefits.....	4
Short-term versus long-term benefits.....	5
Four categories of beneficiary companies.....	5
Four categories of at-risk companies.....	6
Paths to monetization.....	6
Margin implications.....	7
Competition.....	8
Open source vs. proprietary.....	8
Vector database: Why they are necessary and debate on feature versus platform.....	9
Overcoming the hurdles to enterprise adoption (costs, hallucinations, data privacy).....	9
Societal implications.....	10
Generative AI vs. Chat.....	12
RBC AI Index.....	13
RBC AI At-Risk Index.....	13
Index of Prior Notes.....	14



## Why GenAI is so important

**We believe generative AI is a seismic change in the technological landscape.** We view this as the fourth big technological revolution in the past 40 years and each of those had a seminal moment, where the technology became mainstream. In the revolution of the Internet, Internet became mainstream with the launch of Netscape. In the revolution of the cloud, cloud became mainstream with Salesforce on the application side and Amazon Web Services on the infrastructure side. In the third revolution defined by mobile, mobile became mainstream with the launch of the iPhone. Now, in this fourth revolution, ChatGPT's launch last November brought this generative AI mainstream. In other words, we believe generative AI is likely to have major implications not just within the realm of technology, but society at large.

**Generative AI will lead to the next generation of start-ups.** Much like the introduction of the iPhone led to an entire economy of mobile applications and the rise of AWS created entire enterprise software companies made to be cloud native, we believe ChatGPT and generative AI will be the foundation for the next generation of tech start-ups. We also believe the cloud workloads associated with these new start-ups will be additive to overall cloud growth. We believe the OpenAI's recent announcements of GPTs and the upcoming OpenAI app store will be this moment that leads to an explosion of new applications.

**The importance of generative AI lies in the potential to transform the way we interact with machines and with each other.** As generative AI models become more sophisticated and more widely deployed, they have the potential to replace or augment many of the routine tasks that we perform, freeing us up to focus on more complex and creative work. At the same time, they can also help us to communicate more effectively with each other and with the machines that surround us, making it easier to access information, solve problems, and connect with others across the world.

## Framework for evaluating GenAI benefits

**We cannot overstate the importance of generative AI as a technology, but we believe it is equally important to identify which companies can benefit from adopting the technology and have a clear path to monetizing the technology.** For circumstances where the benefit is likely to be purely volume (e.g., Twilio seeing higher volumes of messaging as companies leverage generative AI conversations over SMS), we believe the benefit could be short term in nature and, candidly, not real generative AI winners, at least not in the same capacity as names like Microsoft, HubSpot, and MongoDB. Our primary concern around short-term volume boost would be the sustainability of the tailwind.

**Is there an AI bubble?** We draw the analogy to the dotcom bubble. Today, we use the Internet in ways that were unimaginable during the peak frenzy of the dotcom bubble in 1999-2000. At the same time, there were a lot of early Internet beneficiaries which ended up failing in the long run, like Webvan, Broadcast.com, Pets.com, and Flooz.com — Bitcoin's grandfather. Similarly, while generative AI use cases are likely to exceed our wildest imaginations, we worry about an AI bubble forming, both in public markets (where AI benefits are overstated or with companies that aren't truly AI companies) and in private markets (we joke adding “.ai” to the end of a company's name will add 10 turns to the multiple). We encourage investors to do thorough due diligence on the technology underpinnings of these companies and differentiate real AI winners from the short-term AI beneficiaries likely to end up short in the long term.

**Differentiating surface-level vs. in-depth GenAI use cases.** As we've mentioned before, we believe monetization of GenAI solutions ultimately comes down to how much value they provide. While some GenAI features will likely be tablestakes, even simple use cases like querying files, drafting emails, or proof-reading writing will be difficult to monetize, in our view, as many of these are cases that can be replicated by having ChatGPT open in a separate browser.

On the other hand, if companies build GenAI solutions that leverage domain expertise and combine multiple data sets (both proprietary and public), then these solutions should generate meaningful time and cost savings, which we believe can result in real monetization

**Finally, we believe large public software companies will begin to acquire smaller AI startups,** partly to accelerate the AI roadmap (effectively outsourced R&D), partly as an acqui-hire (since hiring AI engineers is very difficult right now, based on our due diligence), and partly for investor narrative (witness the positive investor reaction to recent AI M&A by data vendors such as Snowflake and Databricks).

### Short-term versus long-term benefits

**It is also important to identify companies that may benefit in the near term, while being hurt in the long term.** For example, as we think almost every company has to adopt a GenAI strategy (even pizza companies), GSIs may benefit from this rush to invest in GenAI. Beyond that timeframe, however, many software companies are using GenAI to make their platforms easier to implement, use, and customize, presenting a long-term headwind for GSIs and IT services, in our view.

**We also believe there are opportunities from a class of stocks where the market is “wrong.”** For example, we believe NICE has been punished as it is viewed as being hurt by GenAI, but we believe it will be a beneficiary, driven by the 3x uplift from selling AI and NICE's data and relationships. On the other hand, we believe PLTR is being treated as an AI beneficiary, when in reality GenAI could hurt the company in multiple ways, including making PLTR's heavy customization less attractive (we estimate services is ~30% of revenue).

### Four categories of beneficiary companies

**1. Large incumbent vendors that can leverage their data and distribution to take advantage of the technology.** Very importantly, the largest companies have to move with a sense of urgency that may not be natural, given how quickly this technology is evolving. We would highlight how quickly Microsoft is moving, as an example. We caution against companies that believe their size alone is a sustainable competitive advantage.

**2. Vertical software leaders that have not just significant data, but the ability to verticalize GenAI solutions** to accelerate the winner-take-all dynamics of vertical software. One of the key pieces of feedback we hear is LLMs are very powerful, but they're also a blank canvas, and it's hard and expensive to train them to speak the language of a specific industry. For example, a theoretical finance-specific LLM would need to know that DCF stands for discounted cash flow and not data classification framework.

**3. Mid-market challengers that can use GenAI as an innovation engine to narrow the gap with market leaders.** Effectively, if large incumbents are too slow to move, we think mid-market competitors can leverage their nimbleness to narrow the gap. Additionally, we think some mid-market challengers could be particularly well positioned to narrow the gap with legacy enterprise software companies by using GenAI to iterate features and functionality faster.

**4. Companies that enable generative AI for others.** This could be said from a technology perspective more broadly and include companies like Microsoft and Nvidia, but could also include consultancies that help enterprises understand the opportunity.

With this framework in mind, we view **COUR, CWAN, GWRE, HUBS, MDB, MSFT, NICE, WDAY,** and **ZM** as the most likely beneficiaries in our universe.



## Four categories of at-risk companies

**1. Legacy, on-premise software companies that cannot truly embrace GenAI, which are cloud-only services.** Some legacy on-prem companies will be at a technological disadvantage for leveraging GenAI and this includes a lot of on-premise companies that haven't yet embraced the cloud.

**2. Companies that don't fully embrace GenAI and rearchitect around it.** We would draw the parallel to the cloud and companies that claim to be embracing the cloud (e.g. Tableau, in our view), versus those that rearchitected for the cloud.

**3. Analytics companies claiming to be AI.** We believe GenAI will expose those companies, especially to the extent they have benefited from corporate AI mandates. We see several companies masquerading with the tag of being an AI company when they are not really doing AI. In our view, we must scrutinize the strategies and announcements we hear to determine if they are real.

**4. Companies whose base functionality could be potentially replaced by generative AI,** for example, we have concerns around the basic work/task management space. This category may be the most important and also the most controversial, since we must think out years in the future.

Therefore, we view **ASAN, PLTR, SMAR,** and **ZI** as the most at-risk in our universe.

## Paths to monetization

**Ultimately, companies must have a thoughtful monetization strategy to be a winner,** and we see that happening in both direct and indirect manners. While we remain very bullish on the long-term opportunity with GenAI and continue to lean on our framework for evaluating beneficiary and at-risk companies, we caution investors that monetization could take time to show up in numbers. For companies that aren't named Microsoft, we're generally modeling immaterial direct GenAI revenue in CY24, a level of conservatism we believe is supported by our checks and company conversations. Microsoft, owing to its early investment in OpenAI in 2019 (long before ChatGPT became a household name), has a headstart in integrating GenAI throughout its portfolio. Few beyond Microsoft have discussed a meaningful monetization strategy, but we see two primary paths to monetization: direct and indirect.

**Direct monetization** is when companies can charge specifically for GenAI products and features, like Microsoft has done with Office Copilot, GitHub Copilot, and Azure OpenAI Services, or paywall a number of generative AI features behind a paywall or premium tier in order to drive upgrades. There's also the ability to monetize GenAI on a consumption level, meaning that as customers use the capabilities, they pay for what they use (which directly aligns customer value with costs). In some ways, a consumption model could drive the best adoption in the near term, since there are fewer barriers to actually trying the generative AI features. As we highlighted in our recent primer on software and recessions (see [here](#)), we expect consumption models to be more sensitive in a recession, but also to recover faster. If this plays out as we expect, this could shift more companies to a consumption model or lead them to adopt consumption elements in their models. In addition to this, we believe GenAI could accelerate that shift to consumption. For example, if GenAI allows users to send 10x more marketing emails than prior, adding a consumption element allows software companies to capture some economics on exponentially higher productivity.

**Indirect monetization** is when GenAI drives more usage of a platform, therefore uplifting revenue (e.g. MSFT Azure or MDB Atlas, which can benefit from the velocity of application development). Also, though indirect monetization strategies, some companies are using GenAI

as a competitive advantage to try and improve win rates versus competitors, attract new customers, and improve customer retention. This will be challenging to measure, but these indirect revenue effects can be important and meaningful in competitive markets.

**There is no “right” path to monetization.** Monetization is dependent on the product, how people use it, and how differentiated the GenAI solutions are. In addition, we expect companies to use a mix of these — for example, we expect features like semantic search to effectively become a requirement for all software vendors, while other more advanced features may be easier to drive direct monetization. A novel approach we would highlight which includes both elements of direct and indirect monetization is Workday's approach — consumption with limits. Workday announced its long-awaited GenAI features and monetization strategy at its customer conference and analyst day in September. Workday intends to effectively offer GenAI capabilities to all customers, but with consumption limits. Beyond this, users would pay for GenAI features on a consumption model. In our view, this could be a compelling, hybrid approach to monetizing GenAI. By bundling it for free, the company could see much higher adoption and usage, versus hiding the features behind a paywall (as a bonus, this usage can drive further customer feedback and innovation). Beyond the limits, customers would pay the company, but only based on how much they use and, ultimately, how much value they get out of it. We would advise management teams to consider this approach, as we believe it aligns the interests of customers, management teams, and investors.

**On the other hand, we believe there are companies where adopting GenAI is tablestakes,** and as a result, it will be difficult to monetize the technology. For example, BOX and DBX are leveraging GenAI to make their platforms more valuable, but we expect monetization to be difficult, especially with Microsoft's presence in the space.

## Margin implications

**Rising compute cost drives persistent margin compression.** While there is no official number, we estimate GenAI workloads cost around 5x more than traditional cloud workloads today. We believe this number will rapidly come down (OpenAI's recent announcements with GPT4 Turbo claim to reduce costs by up to 2.75x), but even in steady-state, costs will likely remain elevated.

**As a result, we believe GenAI will create a significant drag on software gross margins.** Much as we saw the move from on-premise to the cloud bring software gross margins from 90% to 75%, we believe GenAI, given the intensity of workloads, could bring software gross margins down to 60%. Having said that, much like cloud also brought a 2x-3x revenue uplift (displacing hardware and personnel spend), we believe GenAI could also bring a revenue uplift of 2x or higher (displacing services and personnel spend). Ultimately, we believe this could be accretive to gross profit dollars, while there is potential for FCF margins to remain flattish (driven by greater OpEx efficiency from utilizing GenAI). However, in the near term, we expect GenAI will be negative gross margin as companies aim to drive adoption prior to monetization.

**GenAI drives R&D efficiency, but then what?** Unsurprisingly, development and engineering has been the area with the most immediate and tangible benefits from GenAI. For example, 46% of new code is developed by AI today and GenAI helps developers code 55% faster (source: GitHub). With this backdrop, we see two routes companies can take. The first is to slow down hiring and drive margin expansion. The second is to invest as aggressively, but utilize this greater efficiency to innovate products, features, and functionality faster than before and improve the actual technology. We would prefer to see growth companies adopt the second path, as companies that focus too much on margin expansion while GenAI is evolving potentially risk getting left behind.



## Competition

**Federated GenAI approaches may be ideal.** While there is constant debate about which LLMs are the best or most powerful (we continue to believe GPT4 remains the most powerful), or which company will win the GenAI wars, we believe the future approach to GenAI will be federated in nature, as there is no one-size-fits-all approach to LLMs. We expect companies to utilize a combination of LLMs, depending on the use case and priorities (e.g. using GPT4 when power is the most important and Claude 2 Instant when speed is the priority), as well as costs and controls (where LLaMa has major advantages). We also expect companies to build or train their own LLMs and potentially incorporate domain-specific LLMs. In our view, this means there could be multiple winners.

**Anthropic — the new “anti-OpenAI/anti-Microsoft”.** As we’ve written before, OpenAI and Microsoft have a meaningful advantage and headstart versus competitors and their fortunes are closely tied, in our view. Meanwhile, we’ve seen Anthropic emerge as the anti-Microsoft, as numerous Microsoft competitors have invested in and partly hitched their wagons to Anthropic (e.g. GOOGL, CRM, ZM, AWS).

**Will GenAI become “shadow IT”?** In the early days of the cloud, tools like Box and Dropbox became “shadow IT” (software or solutions that are used by departments without the buy-in or control of central IT), which ultimately became a security risk and led to greater cloud adoption by the IT organization. We believe the same will happen with GenAI solutions. According to a Blackberry survey, 75% of organizations are considering or implementing bans on ChatGPT and other GenAI solutions in the workplace (61% of organizations intend for these measures to be long-term or permanent). However, based on our due diligence, knowledge workers are still using GenAI for work, often on their personal devices. In our view, this creates additional risk, especially when it comes to proprietary or sensitive data, and, similar to the cloud, we could see IT departments reverse course in order to have some control and overview over the use of GenAI systems.

## Open source vs. proprietary

**Are LLMs commoditized?** Ever since the leaked Google memo in April (titled: “We have no moat, and neither does OpenAI”), investors are increasingly concerned around competition from open source generative AI models and whether the models themselves will become a commodity. While some open source models are impressive (especially LLaMa), our experience and due diligence shows there is still a meaningful gap between the most powerful LLMs and the others in the space. We continue to find GPT4 to be the most powerful and least prone to hallucinations and our checks indicate this is the common perception, even with software developers. We are certainly paying attention to competition between the major players in the space, but feel OpenAI has a major advantage for the time being, especially when we start considering mind share.

**We believe proprietary LLMs still have major advantages, including power and scalability.** Additionally, one underappreciated piece is IP indemnification, which protects the customer from potential legal risks, with the vendor taking on these risks — most notably, Microsoft announced a Copilot Copyright Committee for customers. Having said that, we see advantages to open source models, including on-premise/private cloud deployment options and cost savings. GPT4 remains the leading LLM, in our view, but we are impressed at the rate at which some open source LLMs are narrowing the gap, most notably LLaMa (from Meta) and Falcon (from TII in Abu Dhabi). Ultimately, we believe enterprises will adopt a hybrid approach, where certain use cases will require proprietary models, while others will utilize open source. We do note that Microsoft is embracing both open source and proprietary, via its close-knit partnership with OpenAI, while also offering LLaMa 2 on Azure and Windows and expanding its partnership with Databricks.



## Vector database: Why they are necessary and debate on feature versus platform

**Vector databases are critical for GenAI and semantic search**, but the big question is whether vector requires a standalone, specialized vendor or whether it's a subset of larger and broader data platforms (the secondary question is whether vector databases can become commoditized). Based on our conversations, we believe vector database (which is, in many ways, the next evolution of graph databases) will become a feature of larger platforms, similar to what happened to graph database software. We've already seen a number of broader data platforms incorporate vector database technology, including MDB and ESTC, as well as smaller private vendors like DataStax and SingleStore. Our due diligence suggests these solutions should be able to scale technologically and that only the largest use cases (particularly large banks and large tech companies) may require a specialized vendor (e.g. Pinecone, Weaviate). In addition, many companies we've spoken with are defaulting to working with existing data vendors for vector solutions before going to specialist vendors. Finally, we would note the most recent update to GPT4 allows developers to include information outside the model, which may bypass the need to vector databases altogether.

**Why RAG architecture matters.** RAG (Retrieval-Augmented Generation) architecture allows users to retrieve external knowledge bases and train LLMs on that specific knowledge to ensure more accurate responses. RAG works by first retrieving relevant data from a database (often a vector database), based on the prompt. This information then augments the model as additional context before providing a response. Finally, the LLM generates a response to the prompt utilizing a combination of pre-trained knowledge and externally-retrieved data. This is especially important with enterprise use cases, as well as domain-specific use cases, in our view.

## Overcoming the hurdles to enterprise adoption (costs, hallucinations, data privacy)

**We expect to see widespread adoption of GenAI by enterprises.** Already today, 92% of Fortune 500 companies are building on OpenAI products (per the company's recent announcement), but we believe this is just the start. In our view, every single company must have a generative AI strategy, whether it's utilizing a platform like OpenAI, working with Microsoft, leveraging GSIs, or building solutions themselves on open source models. Companies that embrace GenAI will substantially increase their competitive advantage versus those that don't, in our view.

**Enterprises have concerns, however.** The primary concerns we hear tend to be: 1) high costs associated with GenAI; 2) "hallucinations" or a GenAI returning responses that are wrong (but confidently wrong); 3) data privacy and residency concerns; and 4) lack of domain expertise that create high ramps to realizing value. Based on our conversations, we see roadmaps to solving all these major concerns over time.

**GenAI costs should come down over time.** Right now, there are multiple factors driving up costs, including potential GPU shortages, high levels of CapEx investments, and the learning curve associated with GenAI. As noted, we believe GenAI workloads cost around 5x more than traditional cloud workloads. This should come down over time as companies work through the GPU shortage, partly by building their own and partly by being more selective about when to use GPUs (Nvidia A100s may be necessary for training most LLMs, but lesser GPUs or even CPUs can be used for some inferencing use cases and even for training smaller models). Notably, most companies aren't monetizing GenAI yet (and likely won't in the near term), as they focus on driving adoption and use cases. Even those that are likely aren't charging a substantial premium to drive profitability — for example, we believe Microsoft may actually be losing money on GitHub Copilot, for now (and expect the same out of M365 Copilot), while we believe \$20/month may be too low for ChatGPT Plus. Over time, however, we expect monetization to ramp, driving

better margins.

**Hallucinations can be avoided with guardrails and human intervention.** Hallucinations may be the hardest problem to solve – after all, if users can't trust the information, why would they rely on it (ignoring the fact that human beings themselves often make up answers to questions they don't know the answers to). Anyone who has used ChatGPT (or Bard or Claude) can attest to chat systems returning answers that are outright incorrect. Based on our conversations, we believe there are a few key steps towards minimizing this issue. The first is to train the model on domain-specific knowledge and include guardrails to prevent the LLM from answering questions outside of that domain (e.g., a customer support bot for a bank should not be answering questions about the weather or vada pav recipes). The second is to train the LLM to answer questions only when there is a high degree of confidence, otherwise the LLM should ask for more details or context and admit when it doesn't have an answer. The third is to use human intervention, where human beings review common cases of hallucination and retrain the model to avoid providing those wrong answers in the future.

**Local models solve the data privacy concerns.** Data privacy and data residency concerns really have two components. The first is wanting to protect customer data, especially with PII (personal identifiable information). The second is companies believe their data is a competitive advantage (which is true to an extent, although perhaps overstated) and those companies don't want their data to train the central model, thus helping out their competitors (e.g., Novo Nordisk wouldn't want its clinical data training GPT4 and then Eli Lilly benefitting from that). Local models largely solve both concerns, in our view. In some cases, it's using open-source models (e.g. LLaMa 2) and deploying them in a private cloud environment and then utilizing that alongside other LLMs. In others, it's using the APIs from those models and trusting that data won't train the central model (a promise that OpenAI and Microsoft have both reiterated).

**Bringing domain expertise to LLMs.** This is the major opportunity for software companies, in our view. If most LLMs are a blank slate, domain-specific models can help a customer get 70% of the way to the finished product. This could be vertical expertise (which creates an underappreciated opportunity for vertical software vendors), department-specific (e.g., Workday creating an HR LLM, HubSpot creating ChatSPot), or use case specific (e.g., customer service chat bots). In our view, this helps drive adoption and faster time to value for GenAI.

## Societal implications

**Governments and regulators will increasingly scrutinize GenAI.** The rise of generative AI brings about a number of ethical and legal concerns. We believe governments throughout the world will create legislation around generative AI, including for the use of generative AI systems, preventing malicious use, and the use of customer data.

**US executive order and funding for Generative AI.** On October 30, 2023, The White House issued an extensive executive order on "Safe, Secure, and Trustworthy AI". One notable aspect of the executive order regarding the promotion of safety and security with implications for our coverage is the invoking of the Defense Production Act. The Commerce Department will be establishing requirements for companies developing foundational models (requirements set to be finalized within 90 days from the EO announcement date), to "notify the federal government when training the model, and share the results of all red-team safety tests." The administration believes these measures will ensure AI systems are safe, secure, and trustworthy before companies make them public. In our view, the impacts from the executive order are yet to be determined, but company commentary will be worth monitoring as more details emerge (particularly for implications related to costs for maintaining compliance).

**Indemnification is key for big tech.** As we highlighted earlier, one major advantage larger tech

companies have is indemnification -- the idea that the vendors will take on legal liability for potential copyright infringement (assuming the tools are used properly). Microsoft, Adobe, and OpenAI are some of the large tech vendors to make such announcements.

**There will be substantial disruption from AI**, just as there has been from every technological and industrial revolution in the past. This time will be different, however, because this is the first industrial revolution to disrupt white collar workers. Jobs like lawyers, journalists, and software developers will be disrupted by Generative AI. In our view, while we are likely still years away from people being replaced by technology, we could absolutely get there, and we need to be prepared for that.

**GenAI could create a bigger separation between the top tiers and bottom tiers within any industry.** With software developers as an example — the developers who are just following orders and building features that they were told to work on will likely become less valuable, since GenAI can actually write code on its own. On the other hand, the best software developers that are coming up with creative ideas and innovating will actually become more valuable since they can be far more productive and iterate ideas much faster. We believe society has a role in helping people reskill and upskill for the new economy, which is something much easier said than done. Most discussions about AI impacting white collar jobs have traditionally focused around tasks done by those at more entry level positions and work that is already being outsourced. Examples include analyzing x-rays, researching prior legal opinions, and summarizing events. However, with the advent of generative AI, we see room for even more senior positions being impacted — for example, not only can generative AI research prior Supreme Court rulings, but it has the ability to write legal arguments based on those prior rulings.

**Widespread adoption of generative AI brings ethical and legal concerns.** If a developer builds an application using code generated by ChatGPT, who owns that code? If a student submits an essay written by ChatGPT, is it their own work? These are tough questions to answer, and we believe there will be endless debate on these topics. In response to the legal side of the equation, we expect governments around the world to take action around generative AI, potentially restricting its use in certain areas and putting guard rails around the use of personal data. In addition, if someone uses generative AI for nefarious reasons, does the creator of that AI hold any liability for not having enough guardrails around safety and compliance? If a sick person uses ChatGPT for medical advice, does OpenAI bear any liability if the advice is wrong or are the current disclaimers enough (or is ChatGPT compared to a website like WebMD instead)?

**Data security and customer privacy is an important consideration.** Another consideration is around how generative AI could impact data security and customer privacy through its collection of data in training models, as well as how the tool could potentially be leveraged for malicious use (e.g. more sophisticated cybersecurity attacks). We have already seen a number of companies such as Amazon and Verizon warn employees on their use of ChatGPT due to concerns that sensitive data could be exposed. Similarly, the potential for this tool to harvest and leverage customer data brings up customer privacy concerns. From a financial materiality perspective, these issues could potentially open companies up to reputational impacts and costs if they aren't properly managing heightened security risks.

**Resource intensity will be worth monitoring.** In terms of environmental impact, concerns have been raised around the resource consumption (energy, water usage) and emissions needed to fuel, train and utilize Generative AI tools. Some of these impacts can be reduced through making the models/processors more efficient and increasing usage of renewable energy. Additionally, a number of benefits could be realized from the application of this technology. For example, generative AI could be leveraged to aid climate scenario analysis, as well as better manage and measure a company's resource use and carbon footprint in their operations and supply chains.



## Generative AI vs. Chat

**Differentiating ChatGPT from Generative AI.** While ChatGPT is a flagship product and use case for GenAI and many treat the terms synonymously, it is important to note that GenAI is much broader than just ChatGPT. Generative AI refers to any AI system that can generate new content, based on pre-trained models. Importantly, this is not just limited to answers to questions (which is typically what we see out of ChatGPT), but content of all forms. This includes image generation (products like Dall-E 3 by OpenAI, Midjourney, and Stable Diffusion) and code generation (such as GitHub Copilot).

**Is chat the future interface of software?** Readers who are old enough may remember when software moved from CLI (command-line interface) to GUI (graphical user interface), highlighted by the move from MS-DOS to Windows. While there have absolutely been advancements in GUI (including mobile and more consumer-like experiences), GUI has remained the predominant paradigm since then. We believe that, over time, software will be augmented by chat and, in some cases, become chat-first. Utilizing chat as instructions to execute a command can save significant time versus pointing-and-clicking through several windows and manually entering information.





## RBC AI Index

### Exhibit 1 - RBC AI Index

		Trading Multiples									Growth Metrics						Margin Metrics								
	EV	EV/Rev			EV/GP			EV/FCF			Revenue			Gross Profit			Gross Margins			Free Cash Flow			Rule of 40		
Ticker	(\$M)	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24
ADBE	265,187	15.1x	13.7x	12.2x	16.9x	15.3x	13.7x	35.9x	33.6x	29.5x	12%	10%	12%	11%	10%	12%	89%	89%	89%	42%	41%	41%	53.4	50.8	53.4
COUR	2,117	4.0x	3.4x	2.9x	6.3x	6.4x	5.5x	NM	NM	NM	26%	20%	15%	33%	(1%)	17%	64%	52%	53%	(10%)	(0%)	2%	16.2	20.1	17.1
CRWD	44,234	19.7x	14.6x	11.3x	26.0x	18.8x	14.6x	65.5x	47.5x	35.8x	54%	36%	29%	53%	39%	29%	76%	78%	78%	30%	31%	32%	84.5	66.2	60.2
CWAN	3,632	12.0x	9.9x	8.3x	16.0x	13.1x	10.9x	71.6x	45.8x	37.4x	NM	21%	20%	NM	22%	20%	75%	75%	75%	17%	22%	22%	NM	42.8	41.7
DDOG	31,630	18.9x	15.0x	12.3x	23.5x	18.4x	15.1x	94.0x	60.0x	50.8x	63%	26%	22%	68%	27%	22%	80%	81%	81%	20%	25%	24%	82.9	50.7	46.4
DT	13,898	12.7x	10.2x	8.7x	15.1x	12.1x	10.2x	46.9x	42.3x	35.1x	335%	24%	18%	384%	25%	18%	84%	85%	84%	27%	24%	25%	362.2	47.9	43.1
ESTC	7,065	6.9x	5.9x	5.0x	9.2x	7.8x	6.6x	NM	89.1x	40.6x	28%	16%	18%	24%	19%	18%	75%	76%	76%	0%	7%	12%	28.9	23.0	30.7
GTLB	5,907	13.9x	10.6x	8.3x	15.6x	11.7x	9.2x	NM	NM	NM	32%	27%	NM	33%	26%	89%	91%	90%	(20%)	(1%)	3%	NM	30.3	30.8	
GWRE	7,238	8.3x	7.8x	7.0x	17.6x	13.5x	11.5x	NM	90.9x	54.7x	14%	7%	12%	4%	30%	18%	47%	57%	61%	(4%)	9%	13%	9.6	15.5	24.5
HUBS	21,035	12.2x	9.9x	8.3x	14.7x	11.8x	9.8x	NM	75.4x	58.1x	33%	23%	20%	36%	25%	20%	83%	84%	84%	11%	13%	14%	43.7	35.7	34.4
MDB	25,489	19.9x	15.8x	12.8x	26.4x	20.6x	16.7x	NM	NM	NM	47%	26%	23%	52%	28%	24%	75%	76%	77%	(3%)	6%	10%	43.5	32.0	32.8
MSFT	2,652,005	13.0x	11.7x	10.2x	19.1x	16.8x	15.1x	44.2x	40.2x	38.0x	10%	11%	14%	9%	13%	11%	68%	70%	68%	29%	29%	27%	39.8	40.1	41.2
NET	21,931	22.5x	17.0x	13.3x	28.8x	21.9x	17.0x	NM	NM	NM	49%	32%	28%	48%	32%	29%	78%	78%	78%	(5%)	9%	9%	43.8	40.7	37.2
NEWR	5,772	6.5x	5.8x	5.1x	8.8x	7.4x	NM	NM	49.9x	NM	18%	11%	15%	27%	19%	NM	74%	79%	NM	1%	12%	NM	19.2	23.0	NM
NICE	10,601	4.9x	4.5x	4.0x	6.7x	6.2x	5.6x	22.1x	20.5x	17.3x	13%	8%	11%	15%	7%	12%	73%	72%	72%	22%	22%	23%	35.2	30.3	34.5
NOW	124,822	17.2x	14.0x	11.6x	20.9x	17.1x	14.1x	57.3x	46.4x	37.4x	23%	23%	21%	25%	23%	21%	82%	82%	82%	30%	30%	31%	53.0	53.5	51.8
PANW	82,738	13.4x	11.0x	9.3x	18.2x	14.5x	12.2x	30.3x	31.2x	24.4x	27%	22%	19%	26%	26%	19%	74%	76%	76%	44%	35%	38%	71.0	57.1	57.0
PATH	7,948	7.5x	6.2x	5.2x	8.8x	7.2x	6.1x	NM	35.4x	26.6x	19%	20%	19%	17%	21%	18%	86%	86%	86%	0%	18%	20%	18.6	38.1	38.5
PD	1,800	4.9x	4.2x	3.6x	5.7x	4.9x	4.2x	NM	31.0x	19.8x	32%	16%	15%	32%	17%	16%	85%	85%	86%	3%	14%	18%	34.3	29.4	33.7
PEGA	4,068	3.1x	3.0x	2.7x	4.2x	4.0x	3.6x	NM	18.3x	14.6x	9%	4%	8%	8%	4%	10%	74%	74%	76%	0%	16%	19%	8.8	20.5	26.8
SNOW	48,789	23.6x	17.7x	13.6x	34.0x	24.3x	18.6x	99.7x	68.2x	50.0x	69%	33%	30%	68%	40%	31%	70%	73%	73%	24%	26%	27%	93.1	59.3	57.5
VEEV	27,886	12.9x	11.8x	9.9x	17.4x	16.1x	13.3x	36.3x	30.0x	27.1x	16%	10%	19%	15%	8%	21%	74%	73%	74%	36%	39%	36%	52.1	49.1	55.6
WDAY	55,159	8.9x	7.6x	6.5x	11.5x	9.6x	8.2x	43.5x	33.4x	27.0x	21%	16%	17%	20%	20%	17%	77%	79%	80%	20%	23%	24%	41.4	39.2	41.0
Mean		12.3x	10.1x	8.4x	16.1x	13.0x	11.0x	53.9x	46.8x	34.7x	44%	19%	19%	46%	21%	19%	76%	77%	77%	14%	19%	21%	58.8	38.9	40.4
Median		12.7x	10.2x	8.3x	16.0x	13.1x	11.2x	45.5x	42.3x	35.4x	26%	20%	19%	26%	22%	19%	75%	78%	77%	17%	22%	23%	43.5	39.2	39.7

Source: FactSet Consensus, priced 11/8/23 AMC

## RBC AI At-Risk Index

### Exhibit 2 - RBC AI At-Risk Index

		Trading Multiples									Growth Metrics						Margin Metrics								
	EV	EV/Rev			EV/GP			EV/FCF			Revenue			Gross Profit			Gross Margins			Free Cash Flow			Rule of 40		
Ticker	(\$M)	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24	CY22	CY23	CY24
ASAN	4,167	7.6x	6.5x	5.7x	8.5x	7.2x	6.3x	NM	NM	NM	45%	18%	13%	45%	18%	13%	90%	90%	90%	(29%)	(5%)	0%	15.8	12.6	13.7
FFIV	8,845	3.3x	3.2x	3.2x	4.0x	3.8x	3.8x	19.5x	14.1x	11.8x	2%	3%	0%	(0%)	5%	1%	81%	82%	83%	17%	22%	27%	18.3	25.8	26.7
PLTR	40,753	21.4x	18.4x	15.4x	26.6x	22.6x	18.7x	NM	70.4x	67.3x	24%	16%	19%	22%	18%	21%	81%	81%	82%	12%	26%	23%	35.2	42.5	42.0
SMAR	4,939	6.4x	5.2x	4.3x	8.2x	6.4x	5.3x	NM	41.0x	28.4x	39%	24%	20%	38%	29%	21%	78%	81%	82%	1%	13%	15%	40.5	36.8	35.2
ZI	6,057	5.5x	4.9x	4.7x	6.2x	5.5x	5.3x	15.3x	14.9x	13.6x	47%	12%	4%	47%	14%	3%	89%	90%	89%	36%	33%	35%	83.2	45.3	38.6
Mean		8.8x	7.6x	6.7x	10.7x	9.1x	7.9x	17.4x	15.1x	30.3x	31%	15%	11%	30%	17%	12%	84%	85%	85%	7%	18%	20%	38.6	32.6	31.2
Median		6.4x	5.2x	4.7x	8.2x	6.4x	5.3x	17.4x	27.9x	21.0x	39%	16%	13%	38%	18%	13%	81%	82%	83%	12%	22%	23%	35.2	36.8	35.2

Source: FactSet Consensus, priced 11/8/23 AMC



## Index of Prior Notes

### **RBC Imagine™: Software & Internet Implications of Generative AI and ChatGPT – Moats and Boats**

In this deep dive, we discuss the evolution of the Generative AI (GAI) secular mega-trend and the potentially transformative effect it could have on technology, enterprise software & internet broadly, and MSFT and GOOGL more specifically.

### **GenAI: Our Framework for Picking Winners/Losers, Margin Impacts, Monetization, Quantifying Benefits**

In this note, we discuss our thoughts on common investor questions regarding GenAI, including our framework for picking winners and losers (along with finding stocks where we feel the market is "wrong"). Additionally, we outline how we think about margin implications of GenAI (both gross and FCF margin) and monetization potential. We also provide some illuminating stats on page two around the value of GenAI.

### **Jaluria Software Download: Vector Databases; Virtual Conferences; and Federated Approaches to GenAI**

Is vector database a standalone platform or a feature? Virtual conferences don't work, in our view. Federated GenAI approaches may be ideal.

### **Jaluria Software Download: Novel GenAI Monetization; Surface vs. In-Depth GenAI; Open Source LLMs**

Consumption with limits — a novel GenAI approach; differentiating surface-level versus in-depth GenAI use cases; open source versus proprietary LLMs; and is Anthropic the new "anti- OpenAI/ anti-Microsoft"?

### **Jaluria Software Download: GenAI Monetization, Analyst Days, and Quality vs. Optical Reacceleration**

The long road to GenAI monetization; Microsoft has a major headstart in GenAI; analyst days mostly down the middle; and not all reacceleration stories are the same (quality vs. optical reacceleration).

### **Jaluria Software Download: Verticalizing LLMs, GenAI as Shadow IT, and M&A as Outsourced R&D**

Vertical software's big opportunity in GenAI; will GenAI become "shadow IT"; and M&A as outsourced R&D.

### **Jaluria Software Download: 4 Paths of GenAI Monetization, Macro Uncertainty, R&D Spending**

The four paths of GenAI monetization; there is no "right path" to monetization; we are not out of the macro woods yet; and watch the R&D line, but also R&D efficiency.

### **Jaluria Software Download: Margins, GenAI Benefits, R&D Efficiency, and LLM Commoditization**

Margin upside has been impressive so far, but is this overcorrecting, or was there that much "fat" to trim; GenAI is transformational technology, but pure volume stories are boring; GenAI drives efficiency, but then what; and are LLMs commoditized?

### **Jaluria Software Download: Consumption Models; PE M&A; GenAI Bubble; Has Macro Stabilized?**

Software is trending towards consumption models...and GenAI could accelerate that; is mid-cap private equity M&A off the table; GenAI is not overhyped, but that doesn't mean there isn't an AI bubble; and has macro really stabilized?



### **Takeaways from Ai4 2023: Strong Enterprise Appetite for GenAI, but Ramp Takes Time**

After attending sessions and speaking with partners, customers, and attendees, we offer our key takeaways: Demand for GenAI from businesses is very high, Most companies are in the pilot or proof-of-concept stage, GenAI spending is a mix of new budget and replacement of existing budget, here is a rapidly growing ecosystem of GenAI companies, highlighting the innovation in the space, ChatGPT is being used at work, but is largely “shadow IT” for now, GPT4 is still the leader, for now, and Microsoft is already offering partners a private version of ChatGPT.

### **RBC Generative AI Update**

Section 1 — Executive Summary and Key Highlights, Section 2 — Generative AI derivatives for Cloud, Datacenters and Chip Manufacturers, Section 3 — Cloud/Hyperscale Financial Highlights, Section 4 — Recent Perspectives on Generative AI, Section 5 — Semiconductors Highlights, Section 6 — Selected companies benefiting significantly from AI demand.

### **RBC Imagine Preparing for Hyperdrive — Themes that will Define Our New Future**

Just as we did in our [Landmark 2018 Imagine Report](#), in this thought leadership study we have taken a global, cross-sector approach to determining the themes that will define our new future. Challenging our global research teams in their respective areas of expertise and coming together to share ideas in recent months, we identified five key themes investors and executives alike across all industries must collectively understand to prepare for the years ahead. The themes are: 1. The Quest for Immortality, 2. The Individual Revolution, 3. Artificial Intelligence Activated, 4. Hybrid Living, and 5. The Great Balancing Act.



## Companies mentioned

Adobe Inc. (NASDAQ: ADBE US; \$585.31; Outperform)  
Asana Inc (NYSE: ASAN US; \$19.91; Underperform)  
Clearwater Analytics Holdings Inc (NYSE: CWAN US; \$19.20; Outperform)  
Cloudflare, Inc. (NYSE: NET US; \$64.58; Outperform)  
Confluent Inc (NASDAQ: CFLT US; \$17.80; Outperform)  
Coursera Inc (NYSE: COUR US; \$18.69; Outperform)  
Coveo Solutions Inc (TSX: CVO CN; C\$10.05; Outperform)  
CrowdStrike Holdings, Inc. (NASDAQ: CRWD US; \$193.83; Outperform)  
Datadog, Inc. (NASDAQ: DDOG US; \$99.99; Outperform)  
Dynatrace, Inc. (NYSE: DT US; \$49.35; Outperform)  
Elastic N.V. (NYSE: ESTC US; \$74.80; Outperform)  
F5, Inc. (NASDAQ: FFIV US; \$155.84; Sector Perform)  
Gitlab Inc (NASDAQ: GTLB US; \$44.74; Outperform)  
Guidewire Software, Inc. (NYSE: GWRE US; \$92.26; Outperform)  
HubSpot, Inc. (NYSE: HUBS US; \$442.15; Outperform)  
Microsoft Corporation (NASDAQ: MSFT US; \$363.20; Outperform)  
MongoDB Inc (NASDAQ: MDB US; \$370.04; Outperform)  
NICE Ltd. (NASDAQ: NICE US; \$169.75; Outperform)  
Pagerduty, Inc. (NYSE: PD US; \$21.18; Outperform)  
Palantir Technologies Inc. (NYSE: PLTR US; \$18.49; Underperform)  
Palo Alto Networks, Inc. (NYSE: PANW US; \$241.08; Outperform)  
Pegasystems Inc. (NASDAQ: PEGA US; \$45.55; Outperform)  
Pro Medicus Limited (ASX: PME AU; AUD85.58; Sector Perform)  
ServiceNow, Inc. (NYSE: NOW US; \$625.33; Outperform)  
Shopify Inc. (NYSE: SHOP US; \$62.74; Outperform)  
Smartsheet Inc. (NYSE: SMAR US; \$40.28; Sector Perform)  
UiPath Inc (NYSE: PATH US; \$17.15; Sector Perform)  
Veeva Systems Inc. (NYSE: VEEV US; \$194.13; Outperform)  
Workday, Inc. (NASDAQ: WDAY US; \$222.60; Outperform)  
Xero Limited (ASX: XRO AU; AUD99.50; Sector Perform)  
Zoom Video Communications, Inc. (NASDAQ: ZM US; \$61.83; Outperform)  
ZoomInfo Technologies Inc (NASDAQ: ZI US; \$13.01; Sector Perform)

## Required disclosures

### Non-U.S. analyst disclosure

One or more research analysts involved in the preparation of this report (i) may not be registered/qualified as research analysts with the NYSE and/or FINRA and (ii) may not be associated persons of the RBC Capital Markets, LLC and therefore may not be subject to FINRA Rule 2241 restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

### Conflicts disclosures

This product constitutes a compendium report (covers six or more subject companies). As such, RBC Capital Markets chooses to provide specific disclosures for the subject companies by reference. To access conflict of interest and other disclosures for the subject companies, clients should refer to <https://www.rbccm.com/GLDisclosure/PublicWeb/DisclosureLookup.aspx?entityId=1>. These disclosures are also available by sending a written request to RBC Capital Markets Research Publishing, P.O. Box 50, 200 Bay Street, Royal Bank Plaza, 29th Floor, South Tower, Toronto, Ontario M5J 2W7 or an email to [rbcinsight@rbccm.com](mailto:rbcinsight@rbccm.com).

The analyst(s) responsible for preparing this research report received compensation that is based upon various factors, including total revenues of the member companies of RBC Capital Markets and its affiliates, a portion of which are or have been generated by investment banking activities of the member companies of RBC Capital Markets and its affiliates.





With regard to the MAR investment recommendation requirements in relation to relevant securities, a member company of Royal Bank of Canada, together with its affiliates, may have a net long or short financial interest in excess of 0.5% of the total issued share capital of the entities mentioned in the investment recommendation. Information relating to this is available upon request from your RBC investment advisor or institutional salesperson.

## Explanation of RBC Capital Markets Equity rating system

An analyst's 'sector' is the universe of companies for which the analyst provides research coverage. Accordingly, the rating assigned to a particular stock represents solely the analyst's view of how that stock will perform over the next 12 months relative to the analyst's sector average.

### Ratings

**Outperform (O):** Expected to materially outperform sector average over 12 months.

**Sector Perform (SP):** Returns expected to be in line with sector average over 12 months.

**Underperform (U):** Returns expected to be materially below sector average over 12 months.

**Restricted (R):** RBC policy precludes certain types of communications, including an investment recommendation, when RBC is acting as an advisor in certain merger or other strategic transactions and in certain other circumstances.

**Not Rated (NR):** The rating, price targets and estimates have been removed due to applicable legal, regulatory or policy constraints which may include when RBC Capital Markets is acting in an advisory capacity involving the company.

### Risk Rating

The **Speculative** risk rating reflects a security's lower level of financial or operating predictability, illiquid share trading volumes, high balance sheet leverage, or limited operating history that result in a higher expectation of financial and/or stock price volatility.

## Distribution of ratings

For the purpose of ratings distributions, regulatory rules require member firms to assign ratings to one of three rating categories - Buy, Hold/Neutral, or Sell - regardless of a firm's own rating categories. Although RBC Capital Markets' ratings of Outperform (O), Sector Perform (SP), and Underperform (U) most closely correspond to Buy, Hold/Neutral and Sell, respectively, the meanings are not the same because our ratings are determined on a relative basis.

Distribution of ratings				
RBC Capital Markets, Equity Research				
As of 30-Sep-2023				
Rating	Count	Percent	Investment Banking	
			Serv./Past 12 Mos.	
			Count	Percent
BUY [Outperform]	820	55.97	250	30.49
HOLD [Sector Perform]	590	40.27	148	25.08
SELL [Underperform]	55	3.75	5	9.09

## Conflicts policy

RBC Capital Markets Policy for Managing Conflicts of Interest in Relation to Investment Research is available from us on request. To access our current policy, clients should refer to

<https://www.rbccm.com/global/file-414164.pdf>

or send a request to RBC Capital Markets Research Publishing, P.O. Box 50, 200 Bay Street, Royal Bank Plaza, 29th Floor, South Tower, Toronto, Ontario M5J 2W7. We reserve the right to amend or supplement this policy at any time.

## Dissemination of research

RBC Capital Markets endeavors to make all reasonable efforts to provide research content simultaneously to all eligible clients, having regard to local time zones in overseas jurisdictions. RBC Capital Markets provides eligible clients with access to Research Reports on the Firm's proprietary INSIGHT website, via email and via third-party vendors. Please contact your investment advisor or institutional salesperson for more information regarding RBC Capital Markets' research.

For a list of all recommendations on the company that were disseminated during the prior 12-month period, please click on the following link: <https://rbcnw.bluematrix.com/sellside/MAR.action>

The 12 month history of Quick Takes can be viewed at [RBC Insight](#).

## Analyst certification

All of the views expressed in this report accurately reflect the personal views of the responsible analyst(s) about any and all of the subject securities or issuers. No part of the compensation of the responsible analyst(s) named herein is, or will be, directly or indirectly, related to the specific recommendations or views expressed by the responsible analyst(s) in this report.

## Third-party disclaimers

The Global Industry Classification Standard ("GICS") was developed by and is the exclusive property and a service mark of MSCI Inc. ("MSCI") and Standard & Poor's Financial Services LLC ("S&P") and is licensed for use by RBC. Neither MSCI, S&P, nor any other party involved in making or compiling the GICS or any GICS classifications makes any express or implied warranties or representations with respect to such standard or classification (or the results to be obtained by the use thereof), and all such parties hereby expressly disclaim all warranties of originality, accuracy, completeness, merchantability and fitness for a particular purpose with respect to any of such standard or classification. Without limiting any of the foregoing, in no event shall MSCI, S&P, any of their affiliates or any third party involved in making or compiling the GICS or any GICS classifications have any liability for any direct, indirect, special, punitive, consequential or any other damages (including lost profits) even if notified of the possibility of such damages.

RBC Capital Markets disclaims all warranties of originality, accuracy, completeness, merchantability or fitness for a particular purpose with respect to any statements made to the media or via social media that are in turn quoted in this report, or otherwise reproduced graphically for informational purposes.

## Disclaimer

RBC Capital Markets is the business name used by certain branches and subsidiaries of the Royal Bank of Canada, including RBC Dominion Securities Inc., RBC Capital Markets, LLC, RBC Europe Limited, RBC Capital Markets (Europe) GmbH, Royal Bank of Canada, Hong Kong Branch, Royal Bank of Canada, Singapore Branch and Royal Bank of Canada, Sydney Branch. The information contained in this report has been compiled by RBC Capital Markets from sources believed to be reliable, but no representation or warranty, express or implied, is made by Royal Bank of Canada, RBC Capital Markets, its affiliates or any other person as to its accuracy, completeness or correctness. All opinions and estimates contained in this report constitute RBC Capital Markets' judgement as of the date of this report, are subject to change without notice and are provided in good faith but without legal responsibility. Nothing in this report constitutes legal, accounting or tax advice or individually tailored investment advice. This material is prepared for general circulation to clients and has been prepared without regard to the individual financial circumstances and objectives of persons who receive it. The investments or services contained in this report may not be suitable for you and it is recommended that you consult an independent investment advisor if you are in doubt about the suitability of such investments or services. This report is not an offer to sell or a solicitation of an offer to buy any securities. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. RBC Capital Markets research analyst compensation is based in part on the overall profitability of RBC Capital Markets, which includes profits attributable to investment banking revenues. Every province in Canada, state in the U.S., and most countries throughout the world have their own laws regulating the types of securities and other investment products which may be offered to their residents, as well as the process for doing so. As a result, the securities discussed in this report may not be eligible for sale in some jurisdictions. RBC Capital Markets may be restricted from publishing research reports, from time to time, due to regulatory restrictions and/or internal compliance policies. If this is the case, the latest published research reports available to clients may not reflect recent material changes in the applicable industry and/or applicable subject companies. RBC Capital Markets research reports are current only as of the date set forth on the research reports. This report is not, and under no circumstances should be construed as, a solicitation to act as securities broker or dealer in any jurisdiction by any person or company that is not legally permitted to carry on the business of a securities broker or dealer in that jurisdiction. To the full extent permitted by law neither RBC Capital Markets nor any of its affiliates, nor any other person, accepts any liability whatsoever for any direct, indirect or consequential loss arising from, or in connection with, any use of this report or the information contained herein. No matter contained in this document may be reproduced or copied by any means without the prior written consent of RBC Capital Markets in each instance.

Additional information is available on request.

### To U.S. Residents:

This publication has been approved by RBC Capital Markets, LLC (member FINRA, NYSE, SIPC), which is a U.S. registered broker-dealer and which accepts responsibility for this report and its dissemination in the United States. Any U.S. recipient of this report that is not a registered broker-dealer or a bank acting in a broker or dealer capacity and that wishes further information regarding, or to effect any transaction in, any of the securities discussed in this report, should contact and place orders with RBC Capital Markets, LLC.

### To Canadian Residents:

This publication has been approved by RBC Dominion Securities Inc. (member CICO). Any Canadian recipient of this report that is not a Designated Institution in Ontario, an Accredited Investor in British Columbia or Alberta or a Sophisticated Purchaser in Quebec (or similar permitted purchaser in any other province) and that wishes further information regarding, or to effect any transaction in, any of the securities discussed in this report should contact and place orders with RBC Dominion Securities Inc., which, without in any way limiting the foregoing, accepts responsibility for this report and its dissemination in Canada.

### To U.K. Residents:

This publication has been approved by RBC Europe Limited ('RBCEL') which is authorized by the Prudential Regulation Authority and regulated by the Financial Conduct Authority ('FCA') and the Prudential Regulation Authority, in connection with its distribution in the United Kingdom. This material is not for general distribution in the United Kingdom to retail clients, as defined under the rules of the FCA. RBCEL accepts responsibility for this report and its dissemination in the United Kingdom.

### To EEA Residents:

This material is distributed in the EU by either RBCEL on an authorised cross-border basis, or by RBC Capital Markets (Europe) GmbH (RBC EG) which is authorised and regulated in Germany by the Bundesanstalt für Finanzdienstleistungsaufsicht (German Federal Financial Supervisory Authority) (BaFin).

### To Persons Receiving This Advice in Australia:

This material has been distributed in Australia by Royal Bank of Canada, Sydney Branch (ABN 86 076 940 880, AFSL No. 246521). This material has been prepared for general circulation and does not take into account the objectives, financial situation or needs of any recipient. Accordingly, any recipient should, before acting on this material, consider the appropriateness of this material having regard to their objectives, financial situation and needs. If this material relates to the acquisition or possible acquisition of a particular financial product, a recipient in Australia should obtain any relevant disclosure document prepared in respect of that product



and consider that document before making any decision about whether to acquire the product. This research report is not for retail investors as defined in section 761G of the Corporations Act.

**To persons receiving this from Royal Bank of Canada, Hong Kong Branch:**

This document is distributed in Hong Kong by Royal Bank of Canada, Hong Kong Branch which is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission. This document is not for distribution in Hong Kong, to investors who are not “professional investors”, as defined in the Securities and Futures Ordinance (Cap. 571 of Hong Kong) and any rules made under that Ordinance. This document has been prepared for general circulation and does not take into account the objectives, financial situation, or needs of any recipient. Past performance is not indicative of future performance. WARNING: The contents of this document have not been reviewed by any regulatory authority in Hong Kong. Investors are advised to exercise caution in relation to the investment. If you are in doubt about any of the contents of this document, you should obtain independent professional advice.

**To persons receiving this from Royal Bank of Canada, Singapore Branch:**

This publication is distributed in Singapore by the Royal Bank of Canada, Singapore Branch, a registered entity licensed by the Monetary Authority of Singapore. This publication is not for distribution in Singapore, to investors who are not “accredited investors” and “institutional investors”, as defined in the Securities and Futures Act 2001 of Singapore. This publication has been prepared for general circulation and does not take into account the objectives, financial situation, or needs of any recipient. You are advised to seek independent advice from a financial adviser before purchasing any product. If you do not obtain independent advice, you should consider whether the product is suitable for you. Past performance is not indicative of future performance. If you have any questions related to this publication, please contact the Royal Bank of Canada, Singapore Branch.

**To Japanese Residents:**

Unless otherwise exempted by Japanese law, this publication is distributed in Japan by or through RBC Capital Markets (Japan) Ltd. which is a Financial Instruments Firm registered with the Kanto Local Financial Bureau (Registered number 203) and a member of the Japan Securities Dealers Association (JSDA) and the Financial Futures Association of Japan (FFAJ).

® Registered trademark of Royal Bank of Canada. RBC Capital Markets is a trademark of Royal Bank of Canada. Used under license.

Copyright © RBC Capital Markets, LLC 2023 - Member SIPC

Copyright © RBC Dominion Securities Inc. 2023 - Member Canadian Investor Protection Fund

Copyright © RBC Europe Limited 2023

Copyright © Royal Bank of Canada 2023

All rights reserved