

CELENT

Solution Brief

BONDGPT

Supporting Fixed Income Analytics and Trading with
Generative AI

Jean Sullivan

September 29, 2023

TABLE OF CONTENTS

<u>EXECUTIVE SUMMARY</u>	<u>3</u>
CONTEXT.....	4
FIXED INCOME ANALYTICS AND TRADING PAIN POINTS.....	6
THE COMPANY	8
<u>KEY BRIEFING TAKEAWAYS.....</u>	<u>10</u>
OVERVIEW OF THE APPLICATION	11
COMPETITIVE DIFFERENTIATION.....	11
<u>DESCRIPTION OF THE APPLICATION</u>	<u>12</u>
BONDGPT ARCHITECTURE AND DATA SOURCES.....	16
HOW LTX AND BONDGPT ADDRESSES TRADER NEEDS AND PAIN POINTS.....	17
HOW LTX ADDRESSED LLM CHALLENGES	19
SOLUTION GAPS.....	21
<u>RESULTS OR MEASURES OF SUCCESS</u>	<u>22</u>
<u>TO CONSIDER.....</u>	<u>23</u>
<u>COPYRIGHT NOTICE.....</u>	<u>25</u>

EXECUTIVE SUMMARY

BondGPT is the first of many generative AI offerings in the Broadridge pipeline. Broadridge has created an innovative and intuitive user experience for bond analysts and traders who use the fixed income analytics and trading platform developed by their subsidiary, LTX. Along with LTX's proprietary AI-based scoring of bond liquidity, similarity, and dealer service, BondGPT addresses many of the pain points felt by its users. Celent believes BondGPT is a powerful tool and looks forward to reporting on its continued evolution.

Context

Generative Artificial Intelligence—the use of algorithms to automatically generate content—leapt into popular culture with OpenAI’s introduction of ChatGPT in November 2022. ChatGPT is one of several large language models (LLMs), which are trained on a corpus of millions of pages of text (and, in many cases, images, videos, and sounds). LLMs are algorithms that, given some input text, can “produce a ‘reasonable continuation’ of text based on ‘what one might expect someone to write after seeing what people have written on billions of webpages, etc.’”¹


For the purposes of this Solution Brief, there are a few key features and challenges that are common among LLMs.


- **Training:** The power of any LLM is a function of how it was trained. For example, to develop an LLM focused on producing résumé cover letters, the LLM would be fed a training set of cover letters. To create the training set, humans could sort through the millions of pages of text available on the internet (with the help of computers, naturally) and label those that are cover letters. By analyzing this training set, the LLM algorithm could determine that the typical cover letter is no longer than a page, includes an introduction, a review of qualifications relevant to the position, and a personal link to the firm’s values and goals, and concludes with a call to action. But presumably the goal is to generate **successful** cover letters—ones that got the writer an interview. Identifying success criteria and feeding the model only “successful” cover letters results in a more powerful and accurate model. In addition to a large training corpus, LLM training typically also includes the use of “Reinforcement Learning from Human Feedback” (RLHF). People hand-generate “correct” responses to queries, which are used to reinforce the LLM’s general training. OpenAI used contractors to create at least 13,000 input/output examples as part of the ChatGPT training set.
- **Temperature:** If an LLM is used to simply generate text by picking the next most likely word to appear within its corpus, the result will often seem flat and devoid of creativity. So, LLMs add a bit of well-chosen “randomness” to produce more interesting output. This randomness is a setting called “temperature,” which ranges from 0 to 1. Settings between 0.0 and 0.4 tend to produce responses that are more focused and consistent. Settings of 0.8 and above generate more diverse and unexpected output. As the OpenAI FAQ notes, “for most factual use cases such as data extraction, and truthful Q&A, the temperature of 0 is best.”
- **Translation:** In addition to generating text, appropriately trained LLMs can excel at language translation. In fact, in 2016 Google began replacing its Google Translate code, which had been developed over the previous decade, with “Transformer” models (the “T” in GPT). In nine months, these models were outperforming the previous models in every language. In addition to text, millions of lines of computer code were also

¹ <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/>

included in ChatGPT’s training corpus. This enables ChatGPT to translate between a human language such as English and a programming language such as Python or SQL. In the example below, GPT-3.5 was used to generate a SQL statement that will extract data about a specific set of bonds from a database. Note that rather than just hard-coding the future maturity date, it was “smart” enough to incorporate the DATEADD and GETDATE functions into the answer.

Default (GPT-3.5)

 write sql code to display all fields in the table BONDS that have a yield greater than 5% and a Maturity greater than 20 years from now

 Sure, here's an example SQL query that retrieves all fields from the `BONDS` table where the yield is greater than 5% and the maturity is greater than 20 years from the current date:

```
sql Copy code
SELECT *
FROM BONDS
WHERE Yield > 0.05
AND Maturity > DATEADD(YEAR, 20, GETDATE());
```

In this query:

- `SELECT *` retrieves all fields from the `BONDS` table.
- `WHERE Yield > 0.05` filters the results to include only rows where the yield is greater than 5%.
- `AND Maturity > DATEADD(YEAR, 20, GETDATE())` adds another filter to include only rows where the maturity date is greater than 20 years from the current date. The `DATEADD` function is used to calculate the future date.

Please note that the exact syntax might vary depending on the specific SQL database system you're using. The above query is a general example and may need adjustments based on your database system's syntax and structure.

- **Tuning:** GPTs are initially trained on a large and diverse corpus, making them suitable for generic tasks like “write a four-paragraph essay about [a given subject]” or “write a valentine in iambic pentameter from me to my wife, whom I first met 61 years ago when we were in preschool.” Since ChatGPT’s corpus included the writings of William Shakespeare, Theodore Geisel (aka Dr. Seuss), and many others, it can also be asked to write in their respective styles, provided the LLM was trained to recognize the authors’ works. OpenAI allows GPT-3.5 Turbo to be fine-tuned by adding to its training set, with fine-tuning for GPT-4 coming in the fall of 2023. GPT is fine-tuned by providing it with both the input prompt text and the ideal generated output. This is likely how OpenAI trained it to extract ingredients and quantities from a recipe to build a shopping list. GPT can be fine-tuned to better “understand” a specific dataset, to follow a set of coding standards, or to output content in a specific style or voice.
- **Questions Involving Math and Numbers:** Since GPT was trained on text, it’s no surprise that it is not particularly good at math, even when the response is a textual answer. Below is an example of how it responded to a question that requires very simple math.



Interestingly, one researcher announced that GPT-3.5 made only two mistakes when presented with the “Math with calculator” section of a 2020 SAT, but later concluded it was likely that SAT questions and the correct answers were included in its training set (probably with RLHF). OpenAI expects to make continual improvements in ChatGPT’s mathematical abilities. However, a Stanford study comparing GPT-3.5 responses with GPT-4 determined that its math proficiency actually decreased dramatically over time.

Fixed Income Analytics and Trading Pain Points

Unlike the equity market, the institutional fixed income market lacks transparency and is a much more complicated environment. Outlined below are the major pain points faced by fixed income traders. BondGPT was developed to resolve these problems.

Information Overload: There are fewer than 8,000 equity securities (stocks and ETFs) traded on US exchanges, plus another 10,000 or so that trade “over the counter.” In contrast, there are over 500,000 unique marketable US corporate bonds, over 1,000,000 marketable municipal bonds from more than 44,000 issuers, and more than 122,000 marketable Treasury securities—plus another 70,000 “stripped” Treasuries. In addition to the same sort of issuer characteristics that equities have (like the sector and industry of their corporate issuer and their ability to pay based on balance sheet, cashflow, and income data), fixed income securities also have tenors, coupon rates, issue dates, maturity dates, call dates, yields, spreads, and ratings. This data often comes from different sources and is stored in different databases across an enterprise. In this environment, the use of traditional database query languages such as SQL requires significant training and resources. Until now, most platforms have had Excel-like user interfaces that allow column-based sorting and filtering.

Time	Ticker	Dlr	Src	Security	B Px	A Px	B Sprd	A Sprd	B Sz	A Sz
09:13	IBM	BBLP	MSG	IBM 2 3/4 01/27/22	101.000	102.000	166.98	104.19		
09:13	TRIPOI	BBLP	MSG	TRIPOI TL 1L USD	89.000	90.500				
09:13	AAPL	BBLP	MSG	AAPL 2.4 05/03/23	99.000	100.000	254.91	218.62		
09:13	PHM	BBLP	MSG	PULTE CDS USD SR 5Y D14	138.000	148.000	138.00	148.00	5MM X 5MM	
				PULTE CDS USD SR 5Y D14	138.000	148.000	138.00	148.00	5MM X 5MM	
				TRIPOI TL 1L USD	89.000	90.500				
				IBM 2 3/4 01/27/22	101.000	102.000	166.98	104.19		
				AAPL 2.4 05/03/23	99.000	100.000	254.91	218.62		
				AAPL 2.4 05/03/23	99.000	100.000	254.91	218.62		
				IBM 2 3/4 01/27/22	101.000	102.000	166.98	104.19		
				PULTE CDS USD SR 5Y D14	138.000	148.000	138.00	148.00	5MM X 5MM	
				TRIPOI TL 1L USD	89.000	90.500				
				IBM 2 3/4 01/27/22	101.000	102.000	166.98	104.19		
				AAPL 2.4 05/03/23	99.000	100.000	254.91	218.62		
				PULTE CDS USD SR 5Y D14	138.000	148.000	138.00	148.00	5MM X 5MM	
				AAPL 2.4 05/03/23	99.000	100.000	254.91	218.62		
				TRIPOI TL 1L USD	89.000	90.500				
				IBM 2 3/4 01/27/22	101.000	102.000	166.98	104.19		
				PULTE CDS USD SR 5Y D14	138.000	148.000	138.00	148.00	5MM X 5MM	

Example: Bloomberg Terminal

Lack of a Single Exchange and Lack of Transparency: Equities are traded on exchanges, and there is complete transparency on prices to ensure that everyone who buys or sells shares know the exchange-wide bid and asked prices. In contrast, there is no single bond exchange. Each broker/dealer has a bond desk that maintains an inventory of bonds that are either owned by the firm or by its customers, and each desk sets the price of the bonds in its inventory. If a trader wants to purchase several of the same bond, they may have to request quotes across multiple desks, which in addition to adding order management complexity can lead to “information leakage,” where a desk learns what bond is being sought, which creates the opportunity to increase pricing. While there is no single “bond exchange,” firms such as LTX enable dealers to create “mini exchanges” by putting their inventories on the platform.

Complex Analytics: Each of the half-million unique corporate bonds is unique based on its issuing company, maturity date, coupon, yield, callability, etc. If one specific bond isn’t available for sale, another very similar bond might be. In addition, the price of each bond is a function of these characteristics plus the current yield curve and the trader’s discount rate and future predictions. Bond analysts and traders are constantly asking, “Is this bond priced appropriately?” and “Is there bond similar to this one with a better price?”

Limited Screen Real Estate: Even with two or more monitors, most bond trading screens are filled with charts and numbers. Traders are loath to give up even a small amount of their screens for yet another application.

The Company

The LTX fixed income trading and analytics platform was launched as a subsidiary of Broadridge in June of 2021. Building off Broadridge's fixed income processing platform, which has 20 of the top 25 fixed income broker dealers as clients, the LTX platform uses its patent-pending LTX AI® and e-protocols to improve liquidity and efficiency for buy side clients and dealers. LTX's mission is to enable AI-driven workflows across liquidity discovery, bond discovery and selection, counterparty selection, and trade execution. The platform is being embraced by leading broker-dealers and asset managers.

Table 1: LTX Firm Profile

YEAR FOUNDED	2020
CUSTOMER BASE	30+ leading broker/dealers 80+ asset managers
GEOGRAPHICAL PRESENCE	North America
EMPLOYEES	60

Source: Celent and LTX

LTX holds several patents and patents-pending for its innovative corporate bond trading technology:

- Patent-pending **bond similarity model** allows users to identify tradeable bonds in the absence of liquidity on a particular bond. The AI-driven model considers a vast range of parameters that allow users to describe how and when they would consider a bond to be similar. The model quickly delivers the most similar bonds based on users' criteria, and the outputs are sortable by similarity and available liquidity on the platform.
- Patent-pending **Cloud Match Scores** allow users to identify and measure real time natural liquidity available on the LTX platform using a neural network.
- Patented **RFX®** trading workflow allows a buy side client to connect with natural buy side liquidity through a dealer, using AI to identify natural counterparties.
- Recently launched data-informed **request for quote (RFQ) workflow with Dealer Selection Scores**. LTX has applied for a patent on the underlying technology of these enhancements to the industry standard RFQ protocol, called RFQ+, including dealer selection analytics and liquidity aggregation.

- All the capabilities and features above are integrated within LTX's OpenAI GPT-4-powered **BondGPT™** application.

KEY BRIEFING TAKEAWAYS

BondGPT represents a new way for fixed income analysts and traders to interact with the mountains of disparate data they deal with every day. It enables them to use natural language to research the market, identify bonds of interest, and trade them on the LTX platform. In conjunction with LTX's proprietary algorithms that score contra interest, find bonds with similar characteristics, and rank bond dealers based in part on their interactions with the user, BondGPT represents the first in what is anticipated to be a series of generative AI-enabled features to be offered throughout Broadridge. Built on the OpenAI GPT chassis, it incorporates "algorithmic guardrails" to ensure accurate results and to comply with SEC guidelines for broker/dealers like LTX.



“Broadridge has been consistently adding real-world AI to boost the value of our offerings. With its ability to provide accurate data and real time insights on bond-related questions, BondGPT has allowed users to experience simplified workflows and improved efficiency. Widely, its use has addressed client pain points and aims to further streamline and automate a market previously stuck in manual processes. We look forward to the evolution of the client experience with further adoption of generative AI throughout our portfolio of products.”

**—Joseph Lo, Head of Enterprise Platforms at
Broadridge**

Overview of the Application

Table 2: Overview of Technology Solution

FINANCIAL INSTITUTION	LTX, a Broadridge Company
INITIATIVE	BondGPT
SYNOPSIS	Powered by OpenAI Large Language Model GPT-4, BondGPT enables fixed income analysts and traders to perform pre-trading analytics of corporate bonds against a broad range of databases using natural language. It is fully integrated with LTX's trading platform as well as with LTX's proprietary AI models that score a bond's liquidity and similarity to other instruments, as well as dealer strength.
TIMELINES	BondGPT was developed over a period of nine months. Though the initial version was being used internally very quickly, much of the subsequent work was to ensure the accuracy of the responses and to install guardrails so its responses comply with SEC guidelines. Separately, the proprietary scoring models were developed over an 18-month period.
TECHNOLOGY USED	<ul style="list-style-type: none"> • Open AI GPT-4 and GPT 3.5 • Whisper (voice recognition) • Proprietary in-house
KEY BENEFITS	<ul style="list-style-type: none"> • Immediate accessibility to information housed across multiple databases. • Intuitive user interface, nominal training required. • Speed of results: Queries take almost no time to create, since they are just natural language, and results generally are returned within 20 seconds. • The chat box is always available but takes up a tiny bit of screen real estate.

Source: Celent and LTX

Competitive Differentiation

While a number of financial services institutions have begun to incorporate generative AI into their products, most are designed to answer a broad range of questions from advisors or their clients. Bank of America has the Erica chatbot for client use and Banker Assist for use by its employees. Morgan Stanley is rolling out a GPT-4-based chatbot for use by its advisors, and Morningstar just announced its Mo chatbot for use by retail customers and advisors. Morningstar claims Mo has “guardrails against entering into certain areas of inquiry.” However, none of these chatbots target fixed income traders or analysts.

What distinguishes BondGPT are its laser focus on the needs of bond analysts and traders, its close integration with LTX's trading platform, and its compliance layer.

DESCRIPTION OF THE APPLICATION

BondGPT query responses generally take one of four forms.

1. The most common response is a filtered table of bonds. For example, the query below returns a list of Apple and Microsoft 30-year bonds with coupons greater than 4%. Basically, BondGPT is translating the natural language query into a database query language, executing the query, and returning the results as a table.

compare APPL 30 year bonds with ticker:MSFT 30 year bonds with coupon greater than 4%

Results for: Comparing 30 year bonds of Apple and Microsoft with coupon rate greater than 4%

Hide Data

Show All Available (▼) entries

Filter within results:

BOND DETAIL	TENOR	ISSUER NAME	S&P RATING	TICKER	COUPON RATE	CUSIP	CLOUD SCORE*	MATURITY DATE	ASSET CLASS	SECTOR	TRADE LEVEL	PRICE	YIELD	SPREAD
APPL 4.85 02/10/53	30yr	Apple Inc.	AA+	APPL	4.85	037833EW6	10	2033-05-10	Investment Grade	TECHNOLOGY	SPREAD	98.43	4.95	71.18
APPL 4.1 08/09/62	30yr	Apple Inc.	AA+	APPL	4.1	037833ER7	10	2062-08-08	Investment Grade	TECHNOLOGY	SPREAD	83.5	5.08	83.64
APPL 4.65 02/23/46	30yr	Apple Inc.	AA+	APPL	4.65	037833BX7	10	2046-02-23	Investment Grade	TECHNOLOGY	SPREAD	95.63	4.97	73.57
APPL 4.25 02/09/47	30yr	Apple Inc.	AA+	APPL	4.25	037833CH1	7	2047-02-09	Investment Grade	TECHNOLOGY	SPREAD	91.18	4.89	64.69
MSFT 4.25 02/09/47	30yr	Microsoft Corporation	AAA	MSFT	4.25	594918CA0	6	2047-02-06	Investment Grade	TECHNOLOGY	SPREAD	92.47	4.79	54.82
MSFT 4.75 11/03/55	30yr	Microsoft Corporation	AAA	MSFT	4.75	594918BM6	5	2055-11-03	Investment Grade	TECHNOLOGY	SPREAD	96.01	5	76.13
MSFT 4.5 02/06/57	30yr	Microsoft Corporation	AAA	MSFT	4.5	594918CB8	5	2057-02-06	Investment Grade	TECHNOLOGY	SPREAD	94.2	4.85	61.31

Note that the query can also contain its own mathematical calculations, so one could request, “List banking & credit bonds with yields higher than the sector average.” Answering these types of questions without BondGPT would likely require a database specialist. Other sample queries include:

- Show me bonds where the bid-ask spread is **tighter** than yesterday, sorted by spread low to high.
- List the top five High Yield Retail bonds trading **above par** based on their **cloud scores**.
- List **Healthcare** bonds trading below \$95 with a yield greater than 6% and amount outstanding for **index inclusion** of \$300 million or more.

Tell more about 857477BA0

Results for: Details of bond with CUSIP 857477BA0

Hide Data

Show All Available (▼) entries

Filter within results:

BOND DETAIL	TENOR	ISSUER NAME	S&P RATING	TICKER	COUPON RATE	CUSIP	CLOUD SCORE*	MATURITY DATE	ASSET CLASS	SECTOR	TRADE LEVEL	PRICE	YIELD	SPREAD
SIT 5.625 PERP H	Perp	State Street Corporation	BBB	SIT	5.62	857477BA0	10	2118-12-15	Investment Grade	FINANCIALS	PRICE	96.25	19.48	N/A

Note the boldfaced “trader jargon” that BondGPT has been trained to understand.

BondGPT maintains the context of conversations, so a user could ask, in sequence:

- a) Which investment grade utility bonds mature between 2025 and 2035?
- b) Show only the 10-year with yields over 6.3%.
- c) Only Gas bonds.
- d) What dealers have this specific bond?
- e) What is its trade activity? (See this final response captured in screenshot below.)

Maintaining context enables users to drill down to focus on subsets of bonds without having to repeat the previous queries, speeding up the search process and avoiding errors that could happen with retyping.

what is its trade activity?

Results for: Trade activity for bond with cusip 86562MAA8

Hide Data

Show entries

Filter within results:

TRACE TRANS DATE	TRACE TRANS TIME	TRACE SIDE	TRACE SIZE	TRACE QUANTITY	TRACE PRICE	TRACE MARKET DATA YIELD	TRACE SPREAD	TRACE REPORTING PARTY TYPE	TRACE CONTRA PARTY TYPE
2023-09-01	00:00:00	B	200000	200000.0	98.92	6.38	89.41	D	C
2023-08-30	00:00:00	S	2.0M	2.0M	98.87	6.42	90.08	D	D
2023-08-30	00:00:00	S	2.0M	2.0M	98.88	6.41	88.83	D	D
2023-08-28	00:00:00	S	1.0M	1.0M	98.77	6.59	101.73	D	D
2023-08-28	00:00:00	S	1.0M	1.0M	98.75	6.59	105.37	D	D

BondGPT can generally determine whether a request should be evaluated in context with a previous one, but just in case, there is a “Reset Query” button that zeroes out the context.

Instead of returning a table of bonds, BondGPT can return summaries, such as:

- Which sectors have the tightest average bid-ask spreads?
- What is the average yield of Verizon bonds?
- Show me average yields by sector for high yield bonds.
- What is the average liquidity score for healthcare investment grade bonds?

- Average options adjusted spread for bonds > BB grouped by sector for the last six months.
2. A table of data detailing or summarizing information from FINRA’s Trade Reporting and Compliance (TRACE) bond trading database.

D2C trades for the month of May grouped and sorted by tenor

Results for: D2C trades for May 2023 grouped and sorted by tenor

Show All Available (▼) entries

Filter within results:

TENOR	TRADE COUNT	TOTAL TRADE VALUE
1Mo	205	2.0B
3Mo	369	5.0B
6Mo	700	18.8B
1Yr	616	9.3B
2Yr	2,779	43.7B
3Yr	2,256	43.7B
5Yr	5,300	141.1B
7Yr	135	10.9B
10Yr	4,000	148.2B
20Yr	1,726	21.8B
30Yr	2,090	80.9B
Perp	359	8.3B

3. BondGPT can answer general questions about fixed income securities, but its guardrails prevent it from providing investment “advice.”

What are US corporate bonds good for?

US corporate bonds have several characteristics:

1. Income Generation: Corporate bonds provide regular interest payments, known as coupon payments.
2. Diversification: US corporate bonds have different risk profiles compared to other asset classes like stocks or government bonds.
3. Capital Preservation: Investment-grade corporate bonds are generally considered less risky than stocks and offer a higher level of capital preservation.
4. Fixed Income Exposure: Corporate bonds offer exposure to the fixed income market, allowing participation in the debt financing of companies.
5. Potential for Capital Appreciation: While corporate bonds are primarily income-generating, there is also the potential for capital appreciation if the bond's market value increases over time.
6. Credit Quality: US corporate bonds are rated by credit rating agencies, providing an assessment of the issuer's creditworthiness.

However, corporate bonds carry risks, including credit risk, interest rate risk, and liquidity risk. It's crucial to understand these risks and one's own risk tolerance when considering any financial instrument.

show me the best bonds

As BondGPT, I am not able to show you the best bonds. However, if you provide specific criteria such as sector, coupon, maturity date, duration, credit rating, price, yield, or any other relevant factors, I can help you find a list of bonds that meet your requirements.

are HY better than IG?

As an AI language model, I don't have personal opinions. High Yield (HY) bonds typically offer higher yields because they are issued by companies with lower credit ratings and higher default risk. These bonds can provide higher returns but also come with higher credit risk.

Investment Grade (IG) bonds, on the other hand, are issued by companies with higher credit ratings and lower default risk. They generally offer lower yields compared to High Yield bonds but are considered more stable and less likely to default.

The choice between HY and IG bonds can be influenced by various factors such as risk appetite and investment goals. Thorough research and consideration of various factors is crucial. Consulting with a financial advisor can also be beneficial.

It can also answer questions like, “What is a Yankee bond?” “What is the Dealer Selection Score?” and “How does RFQ+ differ from RFQ?”

4. Finally, like ChatGPT, BondGPT can generate text based on a prompt, such as the two following examples. These responses automatically go through LTX’s compliance rules—more on that below.

Explain bonds to a middle schooler

Sure! Bonds are a type of financial instrument that individuals can acquire. When one acquires a bond, it essentially means lending money to a company or the government. In return, they commit to repay the money with interest after a certain period of time. It's akin to lending money to a friend who promises to repay with a little extra as a token of gratitude. Bonds are often viewed as safer financial instruments because they have a fixed interest rate and a specific maturity date when the initial amount is returned. It's somewhat similar to a loan made to generate some additional funds.

explain bonds to a finance professor

Bonds are financial instruments that signify a loan made by an investor to a borrower, typically a government or corporation. When an investor acquires a bond, they are essentially lending money to the issuer in exchange for regular interest payments and the return of the principal amount at maturity.

Bonds have several key components:

1. Face Value: The initial amount borrowed by the issuer, which is typically repaid at maturity.
2. Coupon Rate: The fixed interest rate paid to the bondholder, usually expressed as a percentage of the face value.
3. Maturity Date: The date on which the bond reaches its full term and the issuer must repay the face value to the bondholder.
4. Yield: The effective interest rate earned by the bondholder, taking into account the purchase price and coupon payments.

Bonds provide a way for issuers to raise capital for various purposes, such as funding infrastructure projects or expanding business operations. They are considered relatively safer compared to stocks because they offer fixed income and have a predetermined repayment schedule.

Bond prices can fluctuate based on changes in interest rates, credit ratings, and other factors. Higher-rated bonds generally offer lower yields but are considered less risky, while lower-rated bonds may offer higher yields but carry a higher

In addition to the query response, most queries generate both a “Did You Mean...” section that restates other similar queries, as well as a “Potential Follow-Up Questions” section that could be used to learn more or drill down. The user can execute each of these queries by clicking on them.

What do you know about TMT bonds?

Did You Mean...

What are the characteristics of TMT bonds in the Technology sector?

Can you provide an analysis of TMT bonds in the Media & Communications sector?

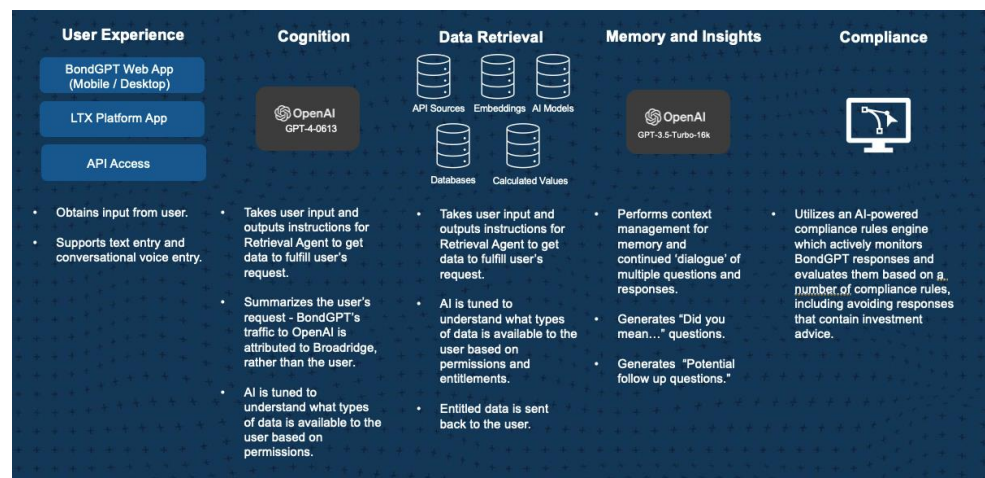
Potential Follow-Up Questions

What is the average yield of TMT bonds in the Technology sector?

How does the total outstanding amount of TMT bonds in the Media & Communications sector compare to other sectors?

In addition, every BondGPT response includes a “Rate This Answer” button, which enables users to provide ongoing feedback in real time. LTX is continuously incorporating this feedback to improve the client experience.

BondGPT Architecture and Data Sources



The graphic above summarizes BondGPT’s base components. GPT-4 is used to translate the user’s natural language query into a language the “retrieval agent” (RA) understands. The RA retrieves the data—from a bond database, trading database, text database, etc.—to fulfill the request. GPT-3.5 (which is much faster than GPT-4) is used to maintain the query context and generate the “Did You Mean...” and “Follow-Up Questions.” Finally, the Compliance Engine filters and edits the results according to proprietary compliance rules.

Table 3: Data Sources

Data	Source	Availability
Bond Details	ICE	Current
Treasury Prices (Intra-Day Benchmark)	ICE	Current
End of Day Corporates Prices	ICE	Current
S&P Credit Ratings	ICE	Current
Bond Similarity model	LTX	Current
Liquidity Cloud information	LTX	Current
Dealer Selection Scores	LTX	Current
Asset manager holdings information (13F)	EDGAR	Current
Public company financial disclosure information (10K, 10Q)	EDGAR	Q4 '23
TRACE EOD	FINRA	Current
Intra-day corporates prices	ICE	Q1 '24
TRACE Intra-day	FINRA	Q1 '24

Source: LTX

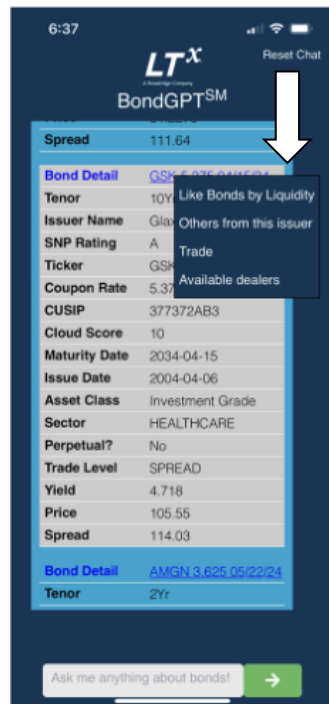
How LTX and BondGPT Addresses Trader Needs and Pain Points

1. Traders Experience Information Overload

Traders do not have to keep track of what data is stored in what database or how to access it. They don't have to learn a separate database query language or request internal IT support. Natural language ensures the user experience is intuitive, and users can even interact with BondGPT by speaking their questions (BondGPT uses Whisper, OpenAI's open source neural network, to translate spoken words).

2. Traders Need Fast Responses

Fixed income traders need to make decisions quickly; bond desk inventory changes constantly, and the bond market responds rapidly to news. BondGPT's intuitive interface has been optimized for speed (which is one reason GPT-4 is used up front but the faster GPT-3.5 is used in the background.) Many responses come back in just a few seconds, with the vast majority completing within 20 seconds. Further, the prewritten "Did You Mean..." and "Follow-Up Questions" are only a click away.



To speed time-to-trade, whenever the details of any bond are displayed, a pull-down menu is available that can:

- a. Show Like Bonds by Liquidity. This uses LTX's proprietary Bond Similarity model, which derives a similarity score based on dozens of bond characteristics, each weighted based on a training set of trader recommendations.
- b. Show other bonds from the same issuer.
- c. **Trade the bond, interfacing directly with LTX's trading system.**
- d. Show the set of available dealers for the bond.

3. Multiple Bond Desks Mean Disjointed Workflows

Without an aggregated exchange such as LTX, traders must communicate with multiple dealers to gather information and execute trades. BondGPT enables liquidity discovery, bond discovery and selection, dealer selection, and trade execution in one interface.

4. Multiple Dealers Leads to a Lack of Pricing Transparency and the Potential for Information Leakage

LTX's Liquidity Cloud allows traders to discover pre-trade liquidity but remain anonymous, ensuring no information leakage. Its Cloud Match score dynamically assesses the real time contra interest in specific bonds.

5. Bond Traders Employ Complex Analytics (which off the shelf LLMs cannot perform)

Bond analysts and traders are constantly asking "is this a good price for this bond?" and "is there a bond that is very similar to this one with a better price?" As noted above, BondGPT users can ask for a list of similar bonds, review their bid/ask, and decide which among them to trade. Similarly, the powerful summarization features make it faster to perform bond analytics.

6. Bond Traders Have Limited Screen Real Estate

The BondGPT chat window can be minimized to take up just an inch or so of screen real estate.

How LTX Addressed LLM Challenges

As previously noted, generic LLMs have several challenges that LTX has directly addressed.

Accuracy: Like humans, LLMs can misunderstand or misconstrue prompts and requests. And, like humans, they don't always realize when they get it wrong. LLMs are designed to respond to prompts regarding virtually any topic (generally within certain guardrails). By reducing the set of questions about which BondGPT can respond to those specific to corporate bonds, LTX significantly reduced the likelihood of it misunderstanding a prompt and delivering an answer to a question that wasn't asked. Many prompt responses include "Did You Mean..." and "Potential Follow-Up Questions" so that users can confirm that BondGPT understood what was being asked.

Lack of Specialized Knowledge: LLMs can answer SAT questions and write college essays; they can write sonnets in the style of Dr. Seuss (or Dr. Dre). But because they have been trained on everything, they lack specialized knowledge. BondGPT's knowledge has been enhanced through Retrieval Augmented Generation (RAG). As IBM explains, "RAG allows LLMs to build on a specialized body of knowledge to answer questions in a more accurate way." This means BondGPT recognizes references to trader jargon such as "junk," "cruise," "beer," or "Yankee bonds" in context, and "understands" bond characteristics like bid/ask spread and ratings.

Hallucinations: One well-documented issue with LLMs is their tendency to "hallucinate"—sometimes they just make things up. This can result in nonsensical answers that often appear perfectly accurate but cite specific events that never actually happened. Hundreds of articles have been written as to why this happens and whether it is a permanent attribute of LLMs. When the original version of BondGPT was asked to display information about Tesla bonds, one of the bonds listed didn't exist (and neither did its CUSIP). LTX mitigates LLM hallucination in two ways. First, through "prompt engineering," it modified the way database queries were made to ensure what is returned are actual values in the database. Second, it alters the temperature parameter depending upon the type of request, so database queries would set the parameter at or near zero but "tell me about Yankee bonds" might have a slightly higher temperature setting.

Dyscalculia: This is the inability to comprehend number-based information and math. BondGPT has been trained to recognize words and phrases such as "average," "mean," "largest," "at least," etc. It understands that an AAA bond has a higher rating than AA, as well as requests for "tight bid/ask spreads."

Compliance: As a subsidiary of Broadridge, LTX is uniquely suited to address compliance issues. This is because Broadridge has for many years offered products that enable advisors to engage with their customers. It has codified its own set of internal compliance rules and implemented them in software. As an example, the ChatGPT team entered this prompt:

"Write a tweet in the style of Taylor Swift explaining why now is the time to invest in the bond market."

Here is what BondGPT generated **without** the Compliance Rule engine's edits:

"Hey Swifties! 🎵 Investing in the bond market is like writing a love story with your money. 💰💖 With interest rates high and stability on the rise, now's the perfect time to dive into bonds. ✍️📧 They offer steady income and protection for your hard-earned cash. So, grab your pen and write your financial future! 📌 #Investing #BondMarket #FinancialLoveStory"

Here is the response after the automated Compliance engine made its edits:

"Hello Swifties! 🎵 The bond market can be seen as a narrative involving your finances. 💰💖 It's characterized by interest rates and stability factors. Bonds are known for their steady income and protection attributes. So, consider this as part of your financial journey! 📌 #BondMarket #FinancialNarrative"

Note the original tweet speaks of now being the “perfect” time and makes predictions about interest rate stability. The post-Compliance version tones these down and suggests “considering” bonds.

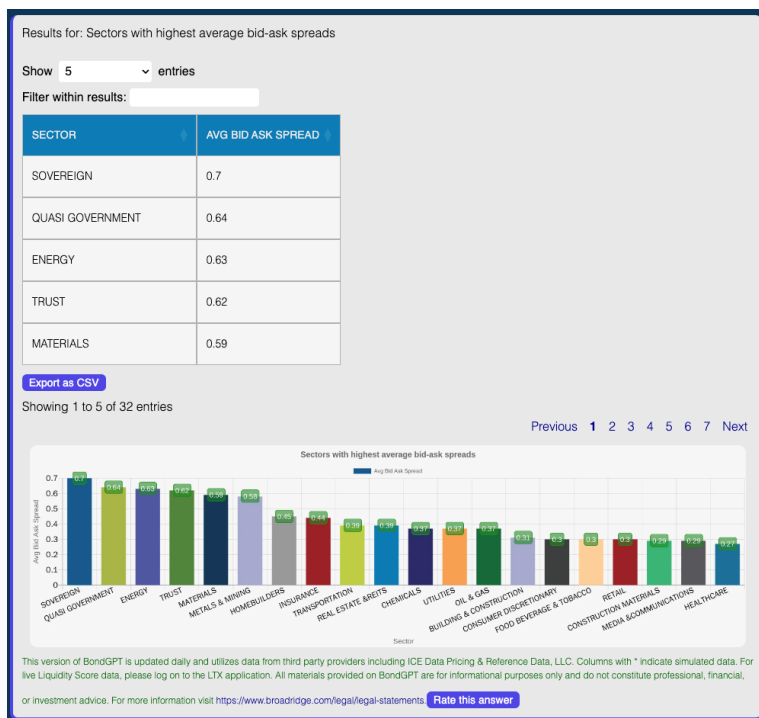
Solution Gaps

- BondGPT doesn't always correctly interpret requests. For example, when asked to display "JPMC bonds maturing between 10 and 20 years from now," it originally thought "now" referred to the ticker symbol for ServiceNow Inc. and just listed a single ServiceNow bond. Fortunately, the "Did You Mean..." made it immediately clear what had happened and rephrasing to "JPMC bonds maturing in 10 to 20 years" returned the expected response.

In another example, a request was made to "display average options adjusted spreads for bonds > BB grouped by sector for the last six months," but the rephrased description of the query it performed was for "bonds rated **below** BB."

When each of these issues were pointed out to LTX, they were resolved in less than 48 hours, but like any LLM-based user interface, there will likely be additional cases where BondGPT does not correctly evaluate the initial query. Fortunately, the rephrasing along with "Did You Mean..." makes it easy to identify these situations. LTX uses ongoing feedback from "Rate This Answer" button to improve response accuracy.

- Many BondGPT responses return a bar chart in addition to a table or summary. LLMs are not normally trained to generate charts, so this is a unique feature of BondGPT which its users have said is quite informative. Each chart has been custom designed based on the data being displayed; however, it is not currently possible to change the graph type (bar versus line, for example) or the axes or the data being displayed. LTX's RoadMap includes enhancements to the charting feature.



RESULTS OR MEASURES OF SUCCESS



“Combining LTX’s AI models with generative AI, BondGPT combs through mountains of raw data and delivers users the insights they need to inform pricing decisions, facilitate security and counterparty selection, and broaden liquidity access. This innovative technology enables users to make use of more data, making pre-trade research and analysis tasks less time-consuming. The feedback has been really positive—since BondGPT’s launch in June, we have seen a steady flow of customer inquiries to learn about the product and how investors can use it to improve their know-how.”

—Jim Kwiatkowski, CEO of LTX

- BondGPT has only been available for a short time, but already it has been embraced by the market: **600+** users across **75+** leading buy and sell side firms.
- LTX has released more than 25 enhancements to BondGPT based on user feedback since launch, with updates ongoing.
- BondGPT has saved up to 75% of scarce front office personnel screen real estate.
- BondGPT provides access to 10M+ data attributes within a tiny chatbox.
- Most questions are answered in less than 20 seconds, which LTX estimates provides time savings of up to 35,900%.
- Simplified front office workflows reduce dependency on several systems to one.
- BondGPT gives decision-makers confidence through verified data on the complete USD-denominated corporate bond market.
- LTX estimates BondGPT offers the ability to save \$2,000/month versus industry market data sources.

TO CONSIDER

Over just the past quarter, dozens of firms, including each of the “Big 4” accounting/consulting firms, have announced billion dollar-plus investments in artificial intelligence, often focusing on generative AI. By focusing its initial offering on a particular set of users—bond traders—with a specific set of pain points that lend themselves to generative AI solutions, Broadridge and LTX have created a truly innovative solution that Celent believes will expand their market share and greatly improve their client’s user experience. The close integration with LTX trading systems and its proprietary scoring algorithms, along with Broadridge’s decades of experience creating compliant client content will make it challenging for competitors to duplicate.

BondGPT is just the first of Broadridge’s generative AI initiatives, and Celent looks forward to the evolution of the next set of these powerful tools.

About Solution Briefs

A Solution Brief (formerly Briefing Note) is a type of Celent Insight launched in 2019 to provide research clients with timely updates on vendor/fintech solutions and strategies. Celent does not charge any fees to write a briefing note, and vendors do not have to be Celent research clients to be eligible for one. However, Celent analysts are selective and publish a limited number of notes throughout a year about briefings they found particularly interesting; the decision whether to write a note is at the Celent analyst's discretion. Vendors have the opportunity to check the draft before it's published to ensure we accurately represent the facts and don't disclose anything confidential, but otherwise do not have editorial control.

RIGHT NOTICE

Copyright 2023 Celent, a division of Oliver Wyman, Inc., which is a wholly owned subsidiary of Marsh & McLennan Companies [NYSE: MMC]. All rights reserved. This report may not be reproduced, copied or redistributed, in whole or in part, in any form or by any means, without the written permission of Celent, a division of Oliver Wyman ("Celent") and Celent accepts no liability whatsoever for the actions of third parties in this respect. Celent and any third party content providers whose content is included in this report are the sole copyright owners of the content in this report. Any third party content in this report has been included by Celent with the permission of the relevant content owner. Any use of this report by any third party is strictly prohibited without a license expressly granted by Celent. Any use of third party content included in this report is strictly prohibited without the express permission of the relevant content owner. This report is not intended for general circulation, nor is it to be used, reproduced, copied, quoted or distributed by third parties for any purpose other than those that may be set forth herein without the prior written permission of Celent. Neither all nor any part of the contents of this report, or any opinions expressed herein, shall be disseminated to the public through advertising media, public relations, news media, sales media, mail, direct transmittal, or any other public means of communications, without the prior written consent of Celent. Any violation of Celent's rights in this report will be enforced to the fullest extent of the law, including the pursuit of monetary damages and injunctive relief in the event of any breach of the foregoing restrictions.

This report is not a substitute for tailored professional advice on how a specific financial institution should execute its strategy. This report is not investment advice and should not be relied on for such advice or as a substitute for consultation with professional accountants, tax, legal or financial advisers. Celent has made every effort to use reliable, up-to-date, and comprehensive information and analysis, but all information is provided without warranty of any kind, express or implied. Information furnished by others, upon which all or portions of this report are based, is believed to be reliable but has not been verified, and no warranty is given as to the accuracy of such information. Public information and industry and statistical data are from sources we deem to be reliable; however, we make no representation as to the accuracy or completeness of such information and have accepted the information without further verification.

Celent disclaims any responsibility to update the information or conclusions in this report. Celent accepts no liability for any loss arising from any action taken or refrained from as a result of information contained in this report or any reports or sources of information referred to herein, or for any consequential, special or similar damages even if advised of the possibility of such damages.

There are no third party beneficiaries with respect to this report, and we accept no liability to any third party. The opinions expressed herein are valid only for the purpose stated herein and as of the date of this report.

No responsibility is taken for changes in market conditions or laws or regulations, and no obligation is assumed to revise this report to reflect changes, events, or conditions that occur subsequent to the date hereof.

For more information, please contact info@celent.com or:

Jean Sullivan

jsullivan@celent.com

Americas

USA

99 High Street, 32nd Floor
Boston, MA 02110-2320

[+1.617.262.3120](tel:+1.617.262.3120)

USA

1166 Avenue of the Americas
New York, NY 10036

[+1.212.345.3960](tel:+1.212.345.3960)

USA

Four Embarcadero Center
Suite 1100
San Francisco, CA 94111

[+1.415.743.7960](tel:+1.415.743.7960)

Brazil

Av. Dr. Chucri Zaidan, 920
Market Place Tower I - 4^o Andar
Sao Paulo SP 04583-905

[+55 11 5501 1100](tel:+55.11.5501.1100)

EMEA

Switzerland

Tessinerplatz 5
Zurich 8027

[+41.44.5533.333](tel:+41.44.5533.333)

France

1 Rue Euler
Paris 75008

[+33 1 45 02 30 00](tel:+33.1.45.02.30.00)

Italy

Galleria San Babila 4B
Milan 20122

[+39.02.305.771](tel:+39.02.305.771)

United Kingdom

55 Baker Street
London W1U 8EW

[+44.20.7333.8333](tel:+44.20.7333.8333)

Asia-Pacific

Japan

The Imperial Hotel Tower, 13th Floor
1-1-1 Uchisaiwai-cho
Chiyoda-ku, Tokyo 100-0011

[+81.3.3500.3023](tel:+81.3.3500.3023)

Hong Kong

Unit 04, 9th Floor
Central Plaza
18 Harbour Road
Wanchai

[+852 2301 7500](tel:+852.2301.7500)