


Solutions Brief

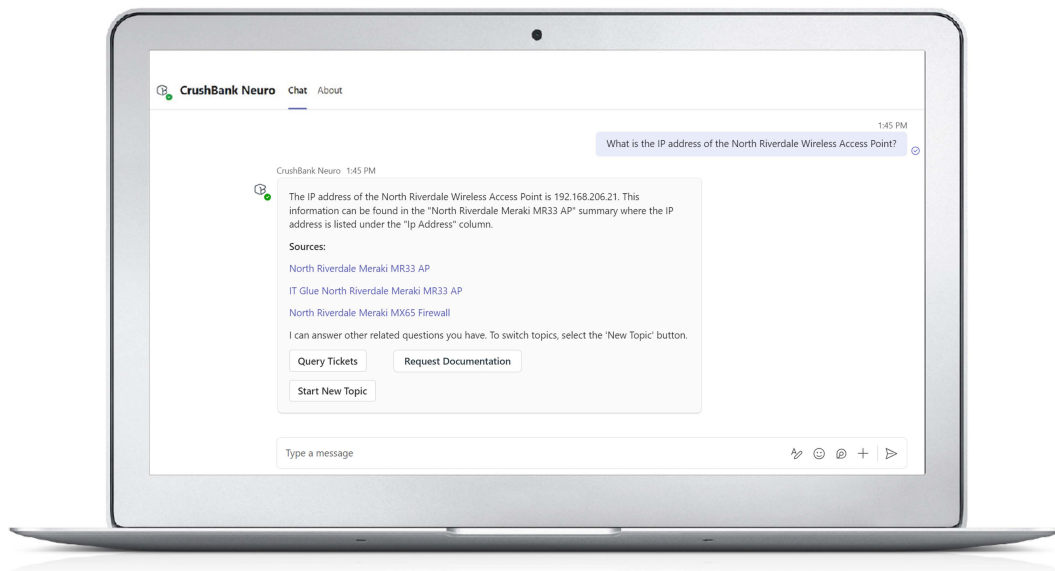
The Power of CrushBank Neuro



Overview

CrushBank is an AI-powered platform designed to help businesses retrieve and leverage critical information from their own data. By utilizing advanced natural language processing (NLP), natural language understanding (NLU), and machine learning, CrushBank Neuro enables organizations to access precise, reliable answers from a wide variety of structured and unstructured data sources.

This Solutions Brief describes how CrushBank works, from data ingestion and normalization to document retrieval and presentation of results.




Data Ingestion and Normalization

One of CrushBank's core capabilities is its ability to ingest data from multiple sources, including both structured data (such as spreadsheets and database reports) and unstructured data (such as notes, manuals, contracts, and documents). These disparate data types often contain critical information needed to solve issues, but they are typically stored in different formats, making them difficult to query or integrate.

- 1. Data Collection:** CrushBank ingests data from various sources that are relevant to answering questions or solving problems. This data could be spread across multiple platforms or databases, and without a centralized system, it can be hard to access efficiently.
- 2. Data Normalization:** During the ingestion process, CrushBank normalizes the data—this means transforming data from its original format into a standardized, consistent structure that can be more easily searched, processed, and stored in the data lake. By ensuring all data is formatted uniformly, CrushBank Neuro enables seamless search and retrieval, regardless of the source or original structure.

For example, data related to network devices, configurations, or support tickets from different tools may have different formats (e.g., different field names or meta data structures). CrushBank standardizes this information, allowing it to be accessed through a unified query process. Additional metadata such as customer names, system types, dates, and times are appended to each piece of data during this phase, further enhancing the system's ability to quickly locate and retrieve precise information when needed.






Enrichment and Vectorization

Once the data is ingested and normalized, CrushBank applies an additional layer of enrichment to further optimize the retrieval process. Enrichment involves adding context to the raw data, enabling Neuro to understand and process more complex queries through meta data such as concepts, categories and sentiment.

1. Enrichment Process: During enrichment, Neuro enhances the data by identifying entities, extracting categories, and recognizing sentiments. For example, for a product specification document, Neuro might extract key information such as the product name, specifications, description and key attributes. These enriched data points are critical in enabling Neuro to perform more accurate searches.

2. Vectorization: Following enrichment, Neuro vectorizes the data. In this process, the enriched data is converted into a mathematical format known as vectors. These vectors capture not only the content but also the context and meaning of the data. This allows Neuro to perform more nuanced searches and comparisons, as the vectorized data provides a deeper understanding of the relationships between terms, phrases, and concepts. By transforming the data into vectors, Neuro enables more advanced searching techniques, optimizing results based on the meaning of a query, not just the keywords used.

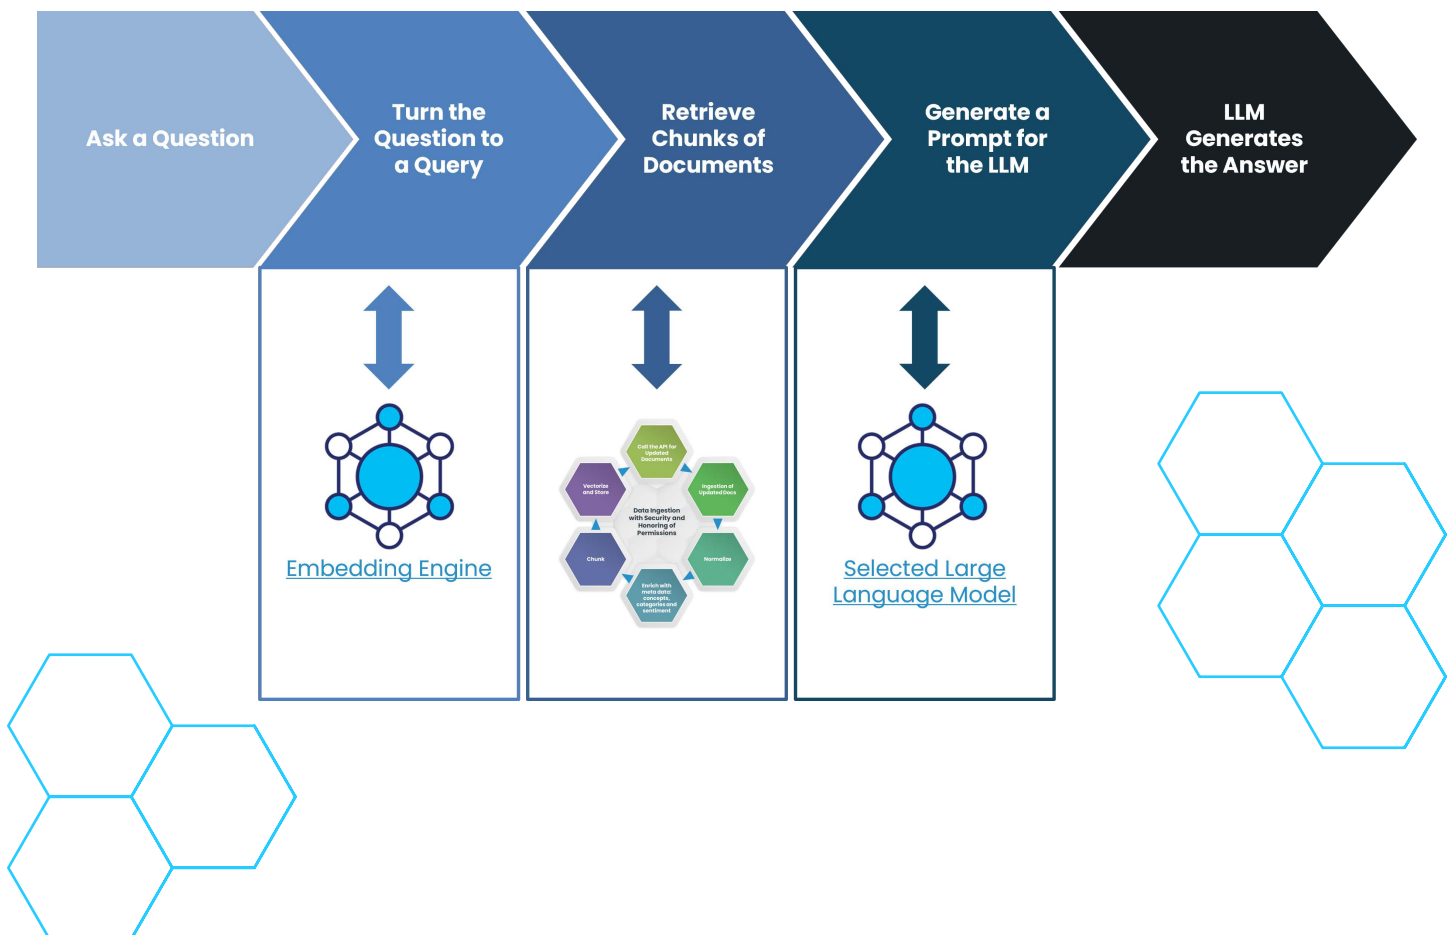
Vectorization allows Neuro to understand context like a client called “Teldar Paper” and the term “firewalls,” determining that the query is related to network infrastructure, even if those exact words aren’t used.



Query Building


When a user poses a question to CrushBank Neuro, the system employs sophisticated query-building techniques to ensure it retrieves the most relevant information from the data lake. This process relies heavily on Neuro's natural language understanding capabilities, which allow it to interpret user questions and convert them into structured, machine-readable queries.

- 1. Language Model Understanding:** Using its integrated language models, Neuro interprets the user's question to understand intent, relevant keywords, and required data types. For example, a question like "How many firewalls does Teldar Paper have?" is parsed to identify key components: the customer (Teldar Paper), the system in question (firewalls), and the type of response required (a numerical count).
- 2. Building the Query:** Neuro then builds a structured query designed to retrieve the most relevant documents or data points. This involves breaking the user's question into specific steps or parts. In the above example, the query would look for information related to the number of firewalls across all of Teldar Paper's documented IT systems. Neuro also takes into account the history of user interactions to refine the query further, ensuring it is as impactful and accurate as possible.



Document Retrieval


Once the query is built, CrushBank Neuro begins the document retrieval process. Neuro searches its data lake, which stores both structured and unstructured data in a normalized and enriched form, for the most relevant documents or document sections.



1. Searching Across Multiple Sources:

Neuro's data lake contains data from multiple sources, including IT service management (ITSM) platforms, Microsoft SharePoint and Teams, customer relationship management (CRM) systems, enterprise resource planning (ERP) systems, and knowledge bases. During the retrieval phase, Neuro can simultaneously search across these diverse sources, identifying both the structured data (e.g., numbers or lists from a database) and unstructured data (e.g., details from technical documents or user manuals).

2. Contextual Retrieval: The data retrieved by Neuro is always contextual. This means that if a piece of information—such as the number of firewalls at a client site—is scattered across different documents, Neuro can combine these chunks into a single, coherent response.




Prompt Generation and Answer

After relevant documents or data sections are retrieved, CrushBank Neuro generates a specific prompt that is sent to the integrated large language model (LLM), to generate a clear and accurate answer. This step ensures that the language model's output remains factual and grounded in the data retrieved from the data lake.

- 1. Creating the Prompt:** The prompt creation process involves packaging the retrieved data into a form that the LLM can easily interpret. Neuro builds boundaries into the prompt to ensure that the answer generated by the LLM remains accurate and does not hallucinate (i.e., provide false or misleading information).
- 2. Generating the Answer:** With the prompt created, the LLM generates a human-readable response. This answer is then double-checked for relevance and accuracy, ensuring that it is fact-based and derived from the retrieved documents. The use of LLMs allows Neuro to provide detailed, precise, and comprehensive answers, even for complex or technical questions.

Presentation of Results

Once Neuro has generated an answer, it is presented to the user in a clear, structured format. Neuro not only provides a concise response but also offers transparency by showing the source documents that were used to generate the answer. This allows users to verify the accuracy of the answer and gain deeper insights if needed. The user can delve deeper with a follow-on question or start a new topic. There is also a button to expand the query results to include tickets. Users are only shown answers that they are entitled to see based on permissions and security.



Prompt Tuning

CrushBank Neuro's prompt tuning ensures that AI responses are highly accurate and relevant, tailored to the specific needs of each organization. By optimizing prompts, Neuro delivers answers that are grounded in the organization's own data and customized to fit a wide range of use cases.

Contextual Precision


Prompt tuning allows Neuro to understand not just the surface-level question but the broader context in which it's asked. Whether the query is related to business operations, technical systems, or customer support, Neuro provides precise, actionable answers that are highly relevant to the specific environment and circumstances.

Grounded, Reliable Responses

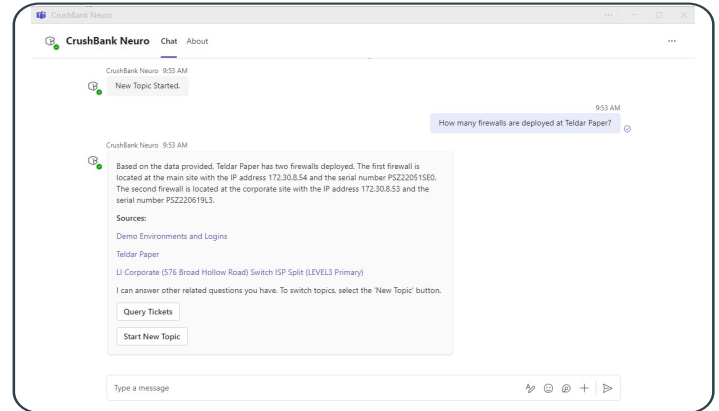
With prompt tuning, Neuro pulls information from trusted sources, including enriched documents and historical data, reducing the risk of speculative or inaccurate answers. This ensures that every response is factual, reliable, and directly linked to the organization's data.

Optimized for Diverse Use Cases

CrushBank's prompt tuning makes Neuro adaptable across various functions, from customer service to operations management. By focusing on the most relevant data and metrics, tuned prompts deliver answers that improve efficiency and decision-making across departments.



Sample Question and Answer:



User Question:


“How many firewalls are deployed at Teldar Paper?”

Step-by-Step Process:

- 1 Reformatting the Query:**
Neuro recognizes that the question involves Teldar Paper and firewalls and translates it into a targeted query designed to retrieve specific numbers related to firewall deployments.
- 2 Query Building:**
The system generates a query that instructs Neuro to search for all documents or records related to firewalls at Teldar Paper, prioritizing records that mention the quantity of deployed systems.
- 3 Document Retrieval:**
Neuro searches through its data lake and retrieves relevant documents, such as network infrastructure reports, configuration management records, or support tickets.
- 4 Prompt Generation:**
The retrieved data is compiled into a prompt for the LLM, instructing it to summarize the information and provide the answer.
- 5 Generated Answer:**
The LLM processes the prompt and generates the following response: “Based on the data provided, Teldar Paper has two firewalls deployed. The first firewall is located at the main site with the IP address 172.30.8.54 and the serial number PSZ22051SE0. The second firewall is located at the corporate site with the IP address 172.30.8.53 and the serial number PSZ220619L3.”

Links are provided to the source documents. A follow-on question can be asked, or a new topic started.

Data Warehouse vs. CrushBank Neuro Data Processing



Traditional data warehouses are limited to structured data and often present challenges when it comes to searching for information across multiple datasets, as each dataset might have a different schema or format. Searching within these systems requires specific queries that must account for the nuances of each dataset.

CrushBank Neuro offers several advantages over traditional data warehouses by:

Handling Both Structured and Unstructured Data

Neuro is designed to integrate both structured data (such as databases or spreadsheets) and unstructured data (such as PDFs or knowledge articles). This allows it to answer questions based on a wider array of information sources.

Normalization and Enrichment

Neuro's data ingestion process includes the normalization of data, which ensures that all data is transformed into a consistent format. Additionally, the enrichment process adds valuable context to the data by tagging it with metadata, such as system names, dates, and entities, allowing for more precise and contextual searches.


Advanced Search Capabilities

By vectorizing the data and utilizing language models, Neuro is able to conduct more intelligent searches, focusing on the meaning and context of queries rather than just keywords. This makes it easier to locate and retrieve the most relevant data, regardless of its original format or source.

Permissions and Access Rights

Neuro respects user access rights and ensures that it only shows information that a user is authorized to see. Data warehouses do not inherently manage access rights on a granular level. Instead, they rely on external security policies or configurations to control what users can access.

The Drawbacks of Creating In-House Solutions with Generic AI



At first glance, it may seem easier to build your own search solution by combining a data warehouse to centralize data with a generic AI model like ChatGPT or CoPilot. While this approach seems promising for solving data search issues, it quickly presents several challenges:

Costs

There are often hidden costs that businesses may not initially anticipate. Many AI platforms charge based on the number of tokens processed, which can be difficult to predict accurately. As usage scales—whether through increased user adoption, more complex queries, or growing data demands—these token-based charges can escalate quickly, leading to unexpectedly high costs. This unpredictability makes budgeting for AI more challenging, especially if a company doesn't monitor or limit token usage effectively. Over time, what starts as a cost-effective solution can turn into a significant financial burden as business needs and AI interactions increase.

Lack of Data Normalization

Centralizing all of your organization's data into a single warehouse is an expensive and time-consuming endeavor. It requires migrating data from various legacy systems, cloud platforms, and proprietary software—each of which operates differently. In the process, data can lose its context, leading to incomplete or fragmented results. Managing these integrations becomes a long-term, costly effort.

Generic AI's Limited Understanding of Your Data

Generic AI models like ChatGPT are designed for broad, conversational applications and aren't equipped to understand the specific needs of an IT support environment. They lack the depth to interpret industry-specific data or the nuances of your internal processes, making their results inconsistent or irrelevant. As a result, using a generic AI can lead to missed information and inaccurate search results.

Security and Governance Concerns

By applying generic AI models to sensitive corporate data, you open the door to significant security and compliance risks. Generic models lack the governance features necessary to ensure data privacy, regulatory compliance, and the level of transparency required by many industries. These risks can lead to data breaches or non-compliant AI usage, resulting in hefty penalties.

Ongoing Maintenance and Scalability Issues

Building and maintaining a custom solution with a data warehouse and generic AI requires ongoing resources. Regular updates, retraining AI models, and scaling to meet growing data demands become continuous hurdles. Over time, the complexity and cost of maintaining this solution can far outweigh the initial investment.

Key Benefits of CrushBank Neuro

Search Across All Data Sources, No Migration Required

Unlike the data warehouse approach, CrushBank Neuro connects directly to your existing data systems—whether they are on-premises, in the cloud, or in legacy systems. This means you don't have to go through costly and time-consuming data migrations. CrushBank Neuro searches across all sources without the need to reorganize your data structure.


AI Powered by Your Data

CrushBank Neuro is built on your own support and client data, not a broad public dataset like those used by generic AI models. This means the search results are always relevant, tailored to your specific industry, and designed to meet your business needs. Whether retrieving support documents, configurations, or customer agreements, CrushBank Neuro delivers results that make sense in the context of your data.

Built on a Strong Foundation with IBM watsonx

CrushBank Neuro leverages the enterprise-grade capabilities of IBM watsonx, which provides a robust architecture designed for secure and scalable data processing. With watsonx, you benefit from:

- **Model Transparency:** Full documentation of how AI models are built and trained, ensuring complete visibility into the data and logic behind your search results.
- **Flexibility and Adaptability:** Choose different models optimized for specific corporate use cases, from HR and finance to legal and customer service.
- **Model Indemnification:** Protect your business from the legal risks associated with AI use through watsonx's built-in model indemnification. This safeguard ensures that your company can deploy AI confidently without worrying about IP infringement claims.
- **Enterprise-Grade Governance:** watsonx offers industry-leading governance capabilities that ensure your data complies with all necessary regulations and security standards. This means your data remains safe, transparent, and auditable—critical for industries with strict compliance requirements.
- **Faster Time to Value:** With CrushBank Neuro, you can start using AI-driven search capabilities almost immediately. There's no need for complex data migrations or lengthy model training periods. CrushBank Neuro integrates seamlessly with your existing systems, reducing the time spent searching for information and increasing overall productivity from day one.



Conclusion:

In conclusion, CrushBank Neuro offers a transformative AI-powered solution tailored to the specific needs of IT support organizations. By seamlessly connecting to both structured and unstructured data, Neuro delivers precise, context-driven answers without the need for costly data migrations. It respects users' access rights. Its use of IBM watsonx provides enterprise-grade security, transparency, and governance, ensuring businesses can rely on accurate, industry-specific results. With its ability to optimize data search and retrieval processes, CrushBank Neuro not only enhances productivity but also mitigates the risks and complexities associated with generic AI solutions. Businesses can confidently leverage Neuro to make informed decisions faster, driving better outcomes across all departments.