



AI Hardware-Accelerated Natural Language to SQL Solution

Transforming Enterprise
Database Access While
Preserving Data Sovereignty

SQL Agent White Paper
May 2025



Abstract

In an era where data accessibility drives business success, organizations face mounting challenges in managing and querying their databases efficiently while maintaining security and compliance. This white paper introduces an innovative solution that combines Neuchips' Viper AI accelerator hardware with GSH's ShareGuru SQLPilot to allow for efficient access to in-house business data such that insight and decision can be made real time.

Note: All performance metrics and capabilities described represent current technology specifications based on internal testing. Results may vary depending on specific implementations and use cases.

AI Hardware-Accelerated Natural Language to SQL Solution

In today's data-driven enterprise environment, organizations face increasing challenges in efficiently accessing and utilizing their database resources. This white paper introduces a revolutionary approach combining Neuchips Viper series specialized AI hardware with GSH's ShareGuru SQLPilot to democratize database access while maintaining enterprise-grade security and performance.



Organizations increasingly rely on database systems for critical operations, yet face significant challenges in today's rapidly evolving data landscape.

The volume of enterprise data continues to grow exponentially, with businesses managing hundreds or even thousands of databases across multiple platforms.



Transforming Database Access

Organizations increasingly rely on database systems for critical operations, yet face significant challenges

Key Impact Metrics

- Reduce database query development time by up to 75%
 - Lower operational costs by reducing dependency on SQL experts
 - Achieve 99.9% data security with on-premises processing
 - Decrease training time for new data analysts by 60%
-
- **Growing complexity of database environments** (average enterprise manages 400+ databases)
 - **Severe shortage of SQL expertise** (demand outpaces supply by 3:1)
 - **Rising database operational costs** (average SQL expert salary: \$120,000+/year)
 - **Increasing demand for rapid data access**
 - **Complex security and compliance requirements**

Feature	Traditional Solutions	Cloud AI Solutions	Our Solution
Processing Location	Local	Cloud	Local
Hardware Required	Multiple Servers	None	Single Card
Data Privacy	Variable	Cloud Risk	Maximum
Power Usage	High	N/A	45W
Integration	Complex	API Only	Flexible
Cost Model	High + Variable	Usage Based	Predictable

Transform hours of SQL development into seconds with hardware-accelerated **AI Agent**. While not instantaneous, our solution dramatically reduces query development time while maintaining accuracy and best practices.

The Database Access Challenge Impact on Business Operations

The traditional approach to database management creates several bottlenecks

- Query development cycles taking 2-3 days on average
- \$150,000+ annual cost per dedicated SQL expert
- 65% of business users waiting for data access
- 40% of queries requiring multiple revisions due to errors



AI Hardware Value Proposition



Hardware Innovation Viper Series LLM Accelerator

The Viper AI Accelerator Card represents a breakthrough in local AI processing, offering comprehensive capabilities for the enterprise AI era.


Dual Purpose Processing

■ LLM Capabilities

- Complete language model processing
- Native BF16 SLM support
- Efficient 45W power consumption
- Flexible thermal solutions

■ Embedded Vector Engine

- Built-in vector similarity search
- Enhanced RAG performance
- Local vector database support
- Real-time embedding processing



Local processing for sensitive data



Your data stays in-house. Process sensitive queries on-premises with zero cloud exposure. Perfect for banking, healthcare, and secure enterprises.

Full **12B** parameter model on a single card



Enterprise AI power, single-card simplicity. Full-scale language model that turns natural language into optimized SQL - **no cloud required.**

Single AI chip Viper card with native BF16 SLM



One chip, endless queries. Native BF16 processing means easier, smarter SQL generation in a single compact card.

45W power efficiency



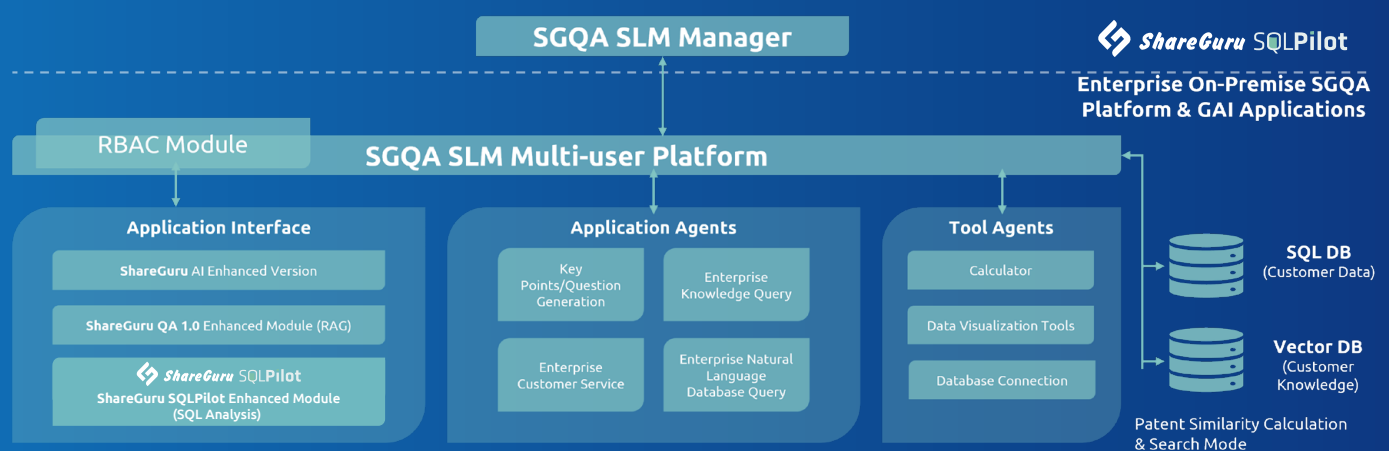
AI at light bulb power. Enterprise-grade query processing at just 45W. Slash your data center costs while going green.

Software Integration

ShareGuru SQLPilot

Solution Offering

- Natural language to SQL conversion with contextual understanding
- Support for complex database operations and multiple query types
- Enterprise-grade security controls
- Compatible with major database platforms



Advanced Features

Query Storage and Retrieval

- RAG (Retrieval Augmented Generation) technology
- Smart similar query lookup
- Consistent response patterns through learned examples

Interactive Refinement

- Natural language field modification
- Conversational query correction
- No programming required for adjustments

Enterprise Integration

- Export capabilities for dashboards
- RESTful API generation
- Automated scheduled analysis
- Step-by-step reasoning mode



Enterprise Data Sovereignty

- Complete control over sensitive data
- No cloud dependency for vector operations
- Eliminate costly pre-training requirements
- Secure on-premises processing
- Custom knowledge base development

Key Features and Benefits

Secure Local Processing

- 100% on-premises data handling
- Zero cloud dependencies
- Sub-second local network latency

Natural Language Query Development

- Average query development time reduced from hours to 3-5 minutes
- 90% reduction in SQL syntax errors
- Built-in query optimization and best practices
- Support for complex analytical queries

Resource Optimization

- 70% reduction in SQL expert dependency
- Operating costs reduced by 45%
- Power consumption under 50W
- Simplified deployment and maintenance

Implementation ROI Analysis

- Typical Enterprise Scenario (250 employees):
- Current SQL Expert Costs: \$450,000/year (3 experts)
- Current Query Development Time: 480 hours/month
- Solution Implementation Cost
- Annual Maintenance
- Projected Annual Savings: \$200,000+
- ROI Achievement: 9-12 months

Conclusion

The Neuchips Viper and Shareguru SQLPilot represents a significant advancement in database accessibility and management.

By combining robust AI hardware with sophisticated natural language processing, we enable organizations to democratize their data access while maintaining security and reducing operational costs.

The solution's ability to run uncompressed 12B parameter models ensures high-quality, accurate query generation, while the 45W power efficiency makes it sustainable and cost-effective for enterprise deployment.

Note: All performance metrics and figures presented are based on internal testing and projections for this emerging technology. Actual results may vary depending on specific implementation and use cases.



Design for AI Inferencing

