

# NeuroBlade SPU

## SQL Processing Unit

Holistic, industry-first, proven acceleration for petabyte-scale analytics processing

## Level Up Your Analytics

**10-100x** performance improvement with **NeuroBlade**

The NeuroBlade SPU revolutionizes data analytics with 10-100x faster SQL query acceleration. The SPU seamlessly integrates with your current infrastructure, tools and queries, offloading SQL tasks from the CPU for optimal efficiency. NeuroLink™ provides a direct, optimized chip-to-chip network that allows for SPUs to expand to multiple nodes without compromising performance.



**SPU G200**

Size	<b>FHHL Form Factor</b>
Power (TDP)	<b>150W</b>
SPU	<b>Up to 4 SPU G200 per server</b>
CPU	<b>--</b>
Memory	<b>--</b>
Storage	<b>Up to 64TB Uncompressed</b>
NeuroLink™	<b>No</b>
Supported OS	<b>Linux</b>



**Dell PowerEdge R760XL**

Size	<b>2U</b>
Power (TDP)	<b>1800W</b>
SPU	<b>Up to 4 SPU G200</b>
CPU	<b>2x Intel Xeon Platinum</b>
Memory	<b>512GB to 4TB</b>
Storage	<b>Up to 256TB Uncompressed</b>
NeuroLink™	<b>No</b>
Supported OS	<b>Linux</b>



**HEQS G100**

Size	<b>6U</b>
Power (TDP)	<b>4000W</b>
SPU	<b>8</b>
CPU	<b>2x Intel Xeon or 2x AMD EPYC</b>
Memory	<b>512GB</b>
Storage	<b>Up to 512 TB Uncompressed</b>
NeuroLink™	<b>Yes</b>
Supported OS	<b>Linux</b>

# NeuroBlade SPU

## SQL Processing Unit


Holistic, industry-first,  
proven acceleration  
for petabyte-  
scale analytics processing

# 10-100x

performance  
improvement  
with Neuroblade

## Illustrative Performance Benchmarks

10-100x  
NeuroBlade SPU



Without SPU



**Hardware-enhanced**  
acceleration at the  
speed of I/O for  
any query engine



**No change to  
existing**  
query platforms  
and file formats  
(e.g. refactoring,  
optimization)



**Consistent  
performance**  
agnostic of query  
complexity



**Accelerate  
AI/ML pipelines**  
for faster and  
more accurate  
results

Illustrative performance of NeuroBlade SPU vs alternatives available, with a reference benchmark running Apache Spark