

## Neo4j on IBM POWER8 – Solutions Data Sheet

### TABLE OF CONTENTS

Introduction	1
Why Graph Databases	1
Why Neo4j	2
Real-Time Graph Processing at Scale	3
IBM Power Systems Built with POWER8	5
Neo4j + POWER8 = The Perfect Solution	6
Conclusion	6

### HIGHLIGHTS

- Neo4j on POWER8 offers 56 TB of extended memory, drastically increasing the size at which real-time graph queries are possible.
- Real-time graph processing with Neo4j on POWER8 supports both standard operational requirements and analytic insights that normally require offline processing.
- IBM POWER8 hardware allows Neo4j to scale both up *and* out for graphs of greater size than ever before.



*Scale your massive graph data like never before with real-time graph processing that's entirely in-memory.*

### Introduction

Graph-based applications and data processing have increased exponentially in the last few years as enterprises realize the transformative nature of the intelligent applications made possible by leveraging data relationships in real time.

Concurrently, the data stored and queried by these applications has also grown exponentially – pushing the art of the possible for larger-scale in-memory graph processing. Yet, even as the amount of data in graph applications continues to rise, enterprises still require real-time performance, with low-latency queries.

Neo4j® on POWER8® is the result of a joint effort between Neo4j and IBM engineering to provide the world's most scalable graph database platform capable of storing and processing graphs of extremely large size all in-memory – shattering all previous real-time scalability limits.

As the world's leading graph database, Neo4j helps business leaders not only manage larger volumes of data but generate insights from their data relationships, and they apply those insights in real time at the point of touch – a task that requires the ability to process massive volumes of data at scale, which IBM® Power Systems® are designed to handle.

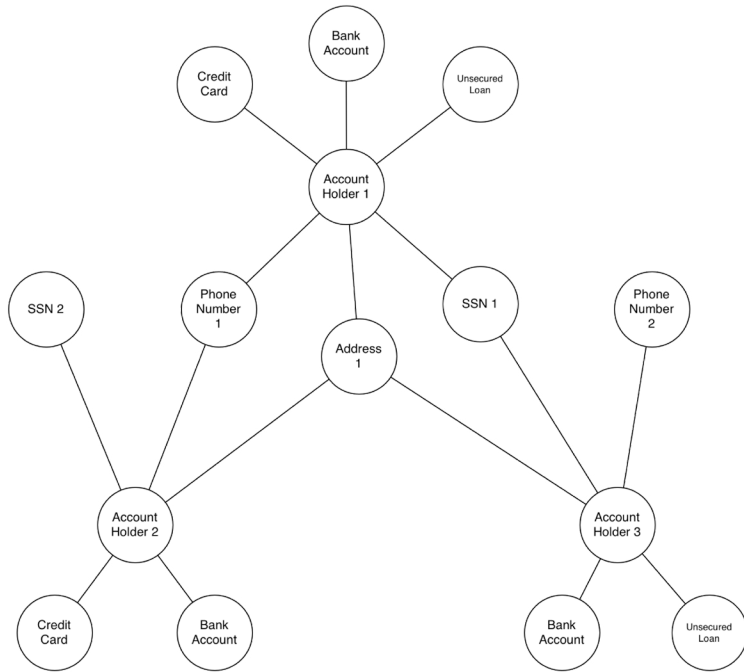
[Neo4j on POWER8](#) offers a number of benefits to today's leading enterprises combining real-time graph processing with the capabilities of IBM hardware to scale up and out like never before.

### Why Graph Databases

Graph databases help enterprise leaders stop merely collecting data points, and start connecting them. By storing data relationship information as a first-class entity, graph databases look beyond individual data points and leverage the relationships that link them.

See Figure 1 below for an example of a basic graph.

## Neo4j on IBM POWER8 – Solutions Data Sheet



**Figure 1. An example of a graph database storing a subset of a potential fraud ring.**

Not only do graph databases effectively store the relationships between data points, but they are also flexible in adding new kinds of relationships or adapting a data model to new enterprise business requirements.

The [top use cases of graph databases](#) include fraud detection, real-time recommendation engines, master data management, network & IT operations, and identity and access management.

Graph databases are taking off as a major enterprise technology. Not only is the volume of data stored in graphs growing bigger, but data models are also becoming increasingly connected.

Within environments of growing data size, data relationships and complexity, enterprises must overcome major challenges of scale in order to retain the real-time advantages of graph insights.

### Why Neo4j

As the clear leader in the graph database space, [Neo4j](#) allows applications to leverage the value of their data relationships. The only limit to how those relationships might be harnessed is the imagination of the user.

With native graph processing and storage, Neo4j delivers real-time, ACID-compliant graph queries, making it perfect for mission-critical applications in an ever-connected world.

Graph databases are the future, and Neo4j helps enterprises step into tomorrow with a significant competitive advantage.

As the clear leader in the graph database space, Neo4j allows applications to leverage the value of their data relationships. The only limit to how those relationships might be harnessed is the imagination of the user.

## Neo4j on IBM POWER8 – Solutions Data Sheet

Only Power Systems provide up to 56 TB of extended memory space with CAPI flash architecture, accommodating graphs of huge size and scale.

### Real-Time Graph Processing at Scale

As a native graph database, Neo4j is set apart by two distinct capabilities that are created specifically for handling graph data: the underlying storage and the processing engine.

Unlike other databases, Neo4j is built from the ground up to store data as a native graph structure to facilitate efficient graph processing. Real-time, native graph processing is another major distinction of Neo4j, supporting both standard operational requirements and bringing real-time analytic insights that would normally (with traditional data processing technologies) require offline processing.

Unlike other technologies that use index lookups to relate data, native graph processing uses memory pointers to relate individual data points: a property often referred to as “index-free adjacency.” With index-free adjacency, queries traverse the graph at an extremely high rate – one million+ pointer chasing operations per second per hardware thread, per query.

The result is “minutes to milliseconds” performance, where queries that take minutes to run on a large RDBMS run in milliseconds on Neo4j.

A key challenge for graph processing at scale is how to handle the sheer size of today's emerging datasets without compromising real-time capabilities. Splitting data up is not an answer, as introducing network hops and index lookups into the middle of a real-time graph query usually means the read (or distributed write) will no longer run in real time.

With the available 56 TB of extended memory on IBM POWER8, the scale of graphs that can be stored and processed in-memory increases drastically, greatly increasing the size at which real-time graph queries are possible.

### IBM Power Systems Built with POWER8

[IBM Power Systems](#) are specifically designed to capture and manage data from a variety of sources and put that data to work in your enterprise – including real-time graph processing.



Only Power Systems provide up to 56 terabytes of extended memory space with CAPI flash architecture, accommodating graphs of huge size and scale.

IBM Power Systems built on POWER8 processors are optimized for a broad range of data workloads, including graph data. IBM Power Systems deliver cost-effective, high-performing and agile IT infrastructure so enterprises can leverage the most advantage from their diverse data sources.

In all big data systems, a balanced system design is key to performance. IBM Power Systems leads the hardware space in three essential areas: processing capability, memory capacity and bandwidth, and input/output (I/O) bandwidth.

The POWER8 is a massively multi-threaded processor. Each core is capable of handling eight hardware threads simultaneously for a total of 96 threads executed simultaneously on a 12-core chip. The processor makes use of large amounts of on- and off-chip eDRAM caches, and on-chip memory controllers enable high bandwidth to memory and system I/O.

## Neo4j on IBM POWER8 – Solutions Data Sheet

POWER8 is capable of clock speeds around 4.35GHz. At this speed, the POWER8 CPU will be around 60% faster than the POWER7 in single-threaded applications and more than twice as fast at multi-threaded tasks.

In addition, the coherent accelerator processor interface – or CAPI – is a new capability that is part of the POWER8 chip. CAPI flash architecture allows off-chip accelerators to have coherent access to POWER8 memory just like traditional power cores. CAPI-attached flash technology introduces a new tier of memory that combines lightning quickness with massive capacity.

Here are just a few more highlights of the POWER8 advantage:

- **Processors:**
  - 4x more hardware threads per core (versus an x86 core)
  - 96 threads on a 12-core chip
  - Up to 1536 threads per system
  - Better Simultaneous Multi-threading (SMT) performance than Hyper-threading (HT)
- **Memory:**
  - 4x more memory bandwidth (versus an x86 core)
  - Up to 16 TB of DRAM
  - Consistent latency as DIMMs are added
  - 192 GBs of sustained memory bandwidth per scale-out socket
  - Up to 230 GBs per second of memory bandwidth for enterprise-class servers
- **Cache:**
  - 4x more cache (versus an x86 core)
  - Up to 231 MB cache per socket (for a 12-core chip)
- **I/O Bandwidth:**
  - 96 GBs per second of peak bandwidth
  - 4x more bandwidth than previous generation (POWER7)

### Neo4j + POWER8 = The Perfect Solution

The combination of Neo4j's native graph processing and storage and POWER8's in-memory vertical scalability is a natural convergence. Neo4j on POWER8 makes it possible to store and process massive-scale graphs in real time – a problem that was simply unsolvable only yesterday.

Neo4j officially supports Linux on IBM Power Systems beginning with Neo4j 2.3. Additionally, Neo Technology and IBM are now accepting beta test participation for POWER8 with CAPI Flash, which will begin Q4 2015. Please [contact us to participate](#).

Together, the strength and performance of Neo4j plus the scalability and speed of POWER8 will provide unrivaled graph application performance.

Neo4j on POWER8 makes it possible to store and process massive-scale graphs in real time – a problem that was simply unsolvable only yesterday.

## Neo4j on IBM POWER8 – Solutions Data Sheet

The union of Neo4j and IBM POWER8 presents a golden opportunity for today's enterprise leaders and application developers.

### Conclusion

The union of Neo4j and IBM POWER8 presents a golden opportunity for today's enterprise leaders and application developers.

With the power of these two technologies, data professionals overcome their big data challenges with a graph database that extracts meaningful, real-time insights from data relationships and the hardware that allows it to scale to massive proportions.

Some enterprise applications that would most benefit from Neo4j on POWER8 include:

- Master datasets connecting organizational, customer and product data together with user activity, inside of a single graph
- Inventory or supply chain management data for global enterprise manufacturers
- Detecting sophisticated fraud across billions of international real-time transactions, whether for insurance, banking or e-commerce
- Monitoring national or international IT networks for outages, dependencies and impact analysis
- IoT applications to manage the vast number of direct and indirect connections between devices, users, locations and transactions

Of course, these are only a few of the many and varied use cases employed by enterprises who harness the strength of Neo4j on POWER8.

Not only does Neo4j on POWER8 deliver real-time graph processing, but the robustness of IBM hardware allows enterprises to scale their graph data to levels never previously imagined.

Neo Technology exists to help the world make sense of data. That's why Neo Technology created Neo4j, the world's leading graph database. Organizations worldwide use Neo4j to extract real-time insights not only from their data but also data relationships. Neo Technology has brought the power of graph databases to organizations both large and small – from enterprises like Walmart, eBay, UBS, Cisco, HP, Telenor and Lufthansa to young startups like Medium and Zephyr Health. These businesses have used Neo4j to build solutions as varied as personalized recommendation engines, identity and access management models, social networks, network monitoring tools and master data management applications. From its inception, Neo Technology has always valued relationships, and now organizations worldwide use the Neo4j graph database to unlock the business value from their data relationships.

For more information on Neo4j,  
contact us via email or phone:  
**1-855-636-4532**  
**[info@neotechnology.com](mailto:info@neotechnology.com)**