

## REPORT REPRINT

This Impact Report was published by 451 Research as part of our syndicated Market Insight subscription service and subsequently licensed for commercial use by Swarm64.

# Swarm64 works with Intel FPGAs to push acceleration engine for relational databases

**JOHN ABBOTT**

**09 JUN 2017**

The Scalable Data Accelerator card, for use within standard x86 servers, acts as a hardware and software offload engine for existing relational databases, boosting performance and capacity to tackle big-data problems on high-velocity data with real-time constraints.

---

THIS REPORT, LICENSED TO SWARM64, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT, WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2017 451 Research, LLC | [WWW.451RESEARCH.COM](http://WWW.451RESEARCH.COM)

It's been two years since we last spoke with Swarm64, and since then the company has honed and simplified its focus considerably. The Berlin-based startup now positions its FPGA-based Scalable Data Accelerator (SDA) card, for use within standard x86 servers, as an accelerator for existing relational databases. It is claiming a performance boost of up to tenfold. The technology currently works with the open source databases PostgreSQL, MySQL and MariaDB, and Oracle support is on its way.

Swarm64 is using Intel field programmable gate array processors and is working in partnership with Intel. Finding a transparent way to boost the power of traditional structured databases (still a huge market) is potentially a very big deal for Intel, which acquired FPGA-maker Altera for \$16.7bn in 2015. Amazon, which already offers EC2 F1 compute instances using Xilinx FPGAs, is another potential partner.

---

## THE 451 TAKE

The market has been moving Swarm64's way, with hardware accelerators gaining traction within on-premises equipment and in the cloud. In most cases, these implementations are relatively transparent to applications, with the primary (mostly x86) CPU still in overall control. However, these accelerators have mostly targeted network acceleration, and new workloads such as AI, deep learning and neural networks. Efforts to run new databases on GPUs have shown significant performance advantages, but the effort involved is huge. Swarm64 is targeting a more general user base with its simpler storage engine plug-in approach. Its best shot at significant traction will be through major partners such as Intel and Amazon.

---

## CONTEXT

Founded early in 2013 by Eivind Liland, Thomas Richter and Alfonso Martinez de la Torre, and led by CEO Karsten Roenner, Swarm64 now has about 20 staff at its offices in Oslo and Berlin – that's double the number from two years ago. It's taken in seed funding, along with a €7.1m (\$7.96m) series A round in September 2015 from Alliance Ventures, Target Partners and Investinor. Two years ago, the company appeared to be positioning itself as a chip company providing a many-core accelerator, with high-speed storage and scalable real-time database as the primary applications. But even then, it was talking about selling its technology as a hardware-software package for data acceleration. It's decided to put all of its efforts behind accelerating relational databases because, despite all the hype around NoSQL and Hadoop, that's where 90% of revenue from database sales still comes from.

## PRODUCTS

There are various ways of boosting database performance, but they are mostly complex or expensive. Swarm64 uses the storage engine interfaces of open source (and potentially commercial) databases to offload the main CPU and take advantage of FPGA chips. It's effectively running its own storage engine and associated software stack alongside the original storage engine. The advantage is that existing databases can be used with no application changes required. MariaDB, MySQL and PostgreSQL are the first three databases supported, with Oracle on the roadmap (using the data cartridge interface). Swarm64 has opted to use Altera (now Intel) FPGAs – hence the Intel partnership – but it could, in theory, also use others, such as Xilinx.

The Scalable Data Accelerator comes in the form of a PCIe card with an Arria 10 FPGA that slots into standard Xeon servers. The FPGA is used for data-compression/decompression and for offloading parts of the OLAP queries. The claim is that, by using FPGA-powered compression, query preprocessing, and more efficient and optimized columnar data structures, the card and software stack will boost both data injection rates and query times tenfold. Data arriving into the database is ready for immediate querying. The hardware costs a few thousand dollars, and Swarm64 will charge about \$1,000 per month for the software stack.

## STRATEGY

The Swarm64 pitch is that it can provide a means of boosting the performance of a traditional RDBMS so that it can begin to tackle big-data problems on high-velocity data with real-time constraints. It's easy to use (i.e., transparent to applications, and with control and configuration that will not be a problem for existing database admin staff) and eliminates the need to move over to a new architecture. Market inertia means that relational databases and SQL will continue to be the backbone technology for data analytics in the enterprise for some time, even though present iterations aren't capable or fast enough to keep up with the data explosion resulting from the growing number of connected devices. At the recent M|17 MariaDB user conference, SDA was shown running MariaDB on a single server, ingesting over a million records per second and analyzing the data in real time.

Swarm64 is currently running customer trials, although none has reached production use yet. The datasets these customers are working on range from less than a terabyte to several terabytes. One trial has focused on replicating a sharded PostgreSQL cluster with a real-time analytics workload. According to Swarm64, a single SDA-enabled server could take the load of eight servers in that cluster. Another trial aimed to prove SDA's transparency to standard analytical tools, with a PostgreSQL customer linking Excel as an analytics front-end tool out of the box. A third trial has been looking at the ingestion of very large datasets from various sources into their working database. Swarm64 sees opportunities in multiple vertical markets, including banking/finance, government, healthcare, IT/telco, insurance, retail and transportation.

At present, SDA covers three of the top-four databases: Oracle (for release soon), MySQL and PostgreSQL. Third-largest player Microsoft is missing from the list, but with Redmond's agreement it could be done. MariaDB is also supported, MongoDB and others using a storage-engine architecture are likely to follow. Intel's forthcoming Xeon+FPGA package (linked together using the QuickPath interconnect) will be supported, as will the more powerful Altera FPGA family, Stratix 10. Horizontal scaling is also on the road map.

## COMPETITION

Despite interest in Hadoop and NoSQL, they represent only a fraction of the database market as a whole. Oracle, Microsoft and IBM still account for 90% of revenue, with AWS the fastest-growing new entrant. A number of start-ups have been experimenting with running databases on GPU hardware, with impressive performance gains. These companies include BlazingDB, Brytlyt, Kinetica, MapD and SQream, but all of them require a shift to a new database and don't support existing standard RDBM products. Ryft has been working on accelerating data analytics on FPGAs with Ryft One, but it's main area of activity has been stream processing – and again, it's a new database rather than an existing one.

Perhaps closer to Swarm64 is Citus Data, which has built a variant of PostgreSQL for scale-out architectures – but while Citus has the horizontal scaling axis covered, it's not looked at the vertical scaling that Swarm64 has focused on. Another company that is scaling open source databases is Percona, which could be competitive in some situations given its acquisition of Tokutek. Tokutek's TokuDB is a high-performance storage engine for MySQL, MariaDB and Percona Server.

Many companies – with greater or lesser credibility – claim to cover both OLTP and OLAP workloads in a single platform, not least Oracle and SAP. Smaller companies with that message include Edge Intelligence, which is also based on PostgreSQL. LeanXcale and Splice Machine are two more players that offer both analytics and transactions. If businesses are primarily looking for an MPP analytics database, contenders include Pivotal (with Greenplum), Hewlett Packard Enterprise (Vertica), Teradata (Aster) and AWS (Redshift).

## SWOT ANALYSIS

### STRENGTHS

With its storage engine plug-in approach, Swarm64 can significantly accelerate existing relational databases - at least those that use a storage engine or similar interface.

### WEAKNESSES

There does need to be an open source storage engine or (in Oracle's case) an open interface, so the majority of business will probably come from existing open source RDBM customers.

### OPPORTUNITIES

Amazon, which has already armed its cloud infrastructure with FPGAs, is the biggest opportunity, and Azure may be the channel through which Swarm64 can get itself into the Microsoft SQL world.

### THREATS

There are many other approaches to accelerating data management, and all the hype is focused elsewhere on analytics and big data using alternative database architectures.