The Emergent Operational/Informational World

An examination of the emerging importance of HTAP

November 2014

A White Paper by

Dr. Barry Devlin, 9sight Consulting
barry@9sight.com

We are experiencing a seismic shift in the way business works. As business and technology merge in the biz-tech ecosystem, real-time data has moved to the center of upfront analytics and immediate action. The old approach to delivering decision support data through a layered architecture is no longer sufficient. A new combined operational/informational world is emerging.

This paper explains how and why the current layered architecture developed and examines why today's business imperatives of speedy decision making and data driven action taking demand a new approach. This operational/informational environment, also called HTAP (hybrid transaction/analytical processing), needs a combination of in-memory operation and novel database techniques to enable simultaneous read/write and long-read activities on the same data.

We describe NuoDB, a modern, distributed, memory-centric database, composed of processes scaling out over multiple hosts in one or more datacenters, which addresses these operational/informational needs. With reduced contention between reads and writes, and a peer-to-peer networked architecture, this database strongly supports the emergent operational/informational needs of today's most advanced business models.

Sponsored by:
www.nuodb.com

Contents

Copyright © 2014, 9sight Consulting, all rights reserved

A new world is emerging. A world where business users don't have to think about using different systems, depending on whether they wanted to use current or historical data. Where staff don't have to distinguish between running and managing the business. Where IT doesn't have to design and manage complex processes to copy and cleanse data from operational systems to data warehouses and marts for business intelligence (BI). As Louis Armstrong says: *"what a wonderful world"*[1]…

There exist a host of examples that demonstrate the dramatic business changes that take place as current and historical data are combined in a more integrated, closed-loop way. The resulting ***biz-tech ecosystem*** restructures existing processes and even reinvents businesses. In many industries, the success—and even survival—of companies will depend on their ability to bridge the current divide between their operational and informational systems. In my 2013 book, *"Business unIntelligence"*[2], I described how the emergence of this biz-tech ecosystem demands a re-evaluation of the traditional layered data architecture to reintegrate operational and informational processing. In 2014, Gartner coined the term hybrid transaction/analytical processing (HTAP)[3] to describe the same need. Furthermore, they pointed to the central role of in-memory databases in its implementation. This technology is evolving rapidly, and a database solution from NuoDB offers a uniquely appropriate approach.

This paper offers you a brief introduction to this new and intriguing world.

How the worlds of operations and decision making diverged

The more we can organize, find and manage information,
the more effectively we can function in our modern world.

Vint Cerf

While the world of flower-power[4] was blossoming on Haight-Ashbury in the mid- to late '60s, academics and progressive business managers were inventing decision support systems (DSS)[5] based on the then novel business applications. Given the technology of the time—batch processing against tape-based files—the approach was simple: run another set of programs against the same files. By the mid-'70s and early '80s, when the wheeling-dealing world of TV's Dallas Ewings[6] was in full bloom, business applications were moving to disk-based databases. In decision support, multiple stand-alone systems emerged, using data extracted from these operational systems.

The separation of operational and informational processing was driven by business needs and technologies that are largely outmoded today.

At the time, business decision-makers operated on planning cycles of months and longer, preferring to focus beyond the fluctuating daily flow of business events. On the technology front, applications and data were handcrafted for operational performance; every byte and processing cycle was optimized. Running DSS simultaneously against the same data was impossible. It was against this background that one of the longest-lived postulates in IT—the need to separate operational and informational systems—arose.

The term *data warehouse* emerged in the industry in the mid-'80s, and Paul Murphy and I published the first formal architecture[7] in 1988. The architecture addressed a key issue in DSS: the multiplicity of data sets being created in large organizations, leading to inconsistent business decisions and operational problems for IT. The data warehouse was largely normalized, and its data reconciled and cleansed on extraction from often incomplete, inaccurate, and inconsistent operational systems. The data warehouse thus provided a complete, accurate and consistent view of the business. Relational databases (RDBMS), which were first commercialized in that era, were adopted as the primary platform for delivering the data

Copyright © 2014, 9sight Consulting, all rights reserved

warehouse for a mix of structural, practical and commercial reasons. Unfortunately, query performance was slow and a third layer of extracted and optimized data marts was added, as shown in figure 1.

These principles, practices and architecture have persisted through to the present day, despite dramatic changes in business needs and technological capabilities. Beyond the current focus on Big Data and the Internet of Things, perhaps the most important evolution in business needs today relates to the speed and agility of BI and analytics. Today's business managers are, in contrast to those of twenty years ago, often keenly interested in seeing the minute-by-minute fluctuations in business measures. Operational BI, as it emerged in the mid-'90s, was a response to this need. However, the layered data warehouse architecture, while promoting data quality, introduces delays in delivering data to business users. Attempts to address this—from operational data stores (ODS) to trickle feeding the warehouse—have only been partially successful. In addition, the design hinders the business demand for agility in a rapidly evolving market.

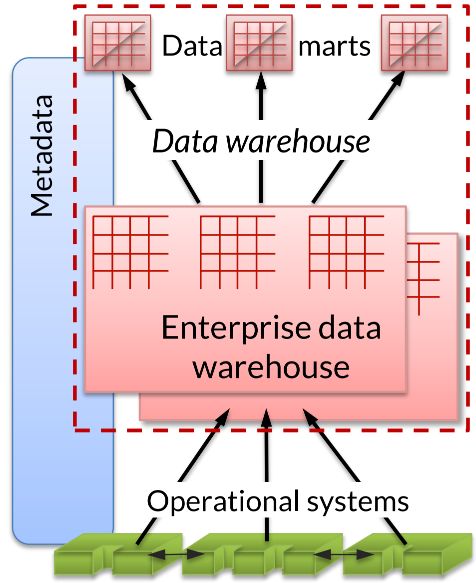

Figure 1: Traditional data warehouse architecture

These limitations were well known, as was the ideal solution—a single data environment. Indeed, a common question I received over the years was if we would ever see a combined operational/informational environment. Until recently, my response has been: *"not in my lifetime"*. Recent technology developments, particularly around in-memory computing, have proven me too pessimistic.

The biz-tech ecosystem drives an integrated world of data

Data is the fabric of the modern world: just like we walk down pavements, so we trace routes through data, and build knowledge and products out of it.

Ben Goldacre

The biz-tech ecosystem is shorthand for a new business world where business needs and technology support exist in tight symbiosis. It's a world that has been emerging for some time now, but the pace of innovation is accelerating so quickly that ignoring it is becoming a life-threatening choice in almost every industry. In the past, business requirements reigned supreme; the role of IT was simply to find the best, cheapest and/or most appropriate solution. Today, technology delivers most or all new business opportunities. It drives the speed of decision making and action taking through ever larger volumes of data and ever more sophisticated analytics. More interestingly, it enables brand new business models where data is often both the raw material and the product.

Of particular interest are the data-related aspects that drive the creation of a combined operational/informational HTAP environment. In terms of business requirements, this equates to an urgent need to analyze data in real-time and to close the analytics-to-action loop with immediate action taking, including the creation of ACID (atomic, consistent, isolated and durable) transactions. Such applications also close the loop between message/event data from social media or the Internet of Things and the process-mediated data managed by traditional business applications.

Let's start with how Google and, indeed, online retailers like eBay get the right content in front of Internet browser-based and mobile app-based consumers at the right time. And in both these environments, the right time basically means instantly.

Copyright © 2014, 9sight Consulting, all rights reserved

Real-time, personalized advertising

Google's business model is based on advertising. When you click on an advertisement, Google notes that as a tiny transaction that it can bill to the advertiser in question. This is, by any definition, Google running their business. It is an operational system that must reliably and correctly record that a particular advert has been clicked and perform the necessary business processes that enable Google to collect its revenue and run its business.

Real-time, personalized advertising provides a prime example of the need for combined operational/informational processing systems.

Of course, Google's particular strength is in knowing what advert might be more likely to interest you at this moment and thus is able to increase the possibility that you might click on it. This is actually an analytical process that occurs in the milliseconds during which the browser or app content is being created before being displayed. All of the information that Google has gathered up to now about you, your preferences and your past behaviors is input to this analysis. Current information, such as your location and details about your current session are also part of the analysis. As are criteria set by the advertiser, data on frequency of display of adverts, and so on. While some of the personal profile information may be analyzed in advance, a significant proportion of the analysis must occur in real-time.

In a traditional BI approach, all analytic work occurs in a different server environment than the operational and transactional work that leads to billing. That separation of concerns necessitates copying data from one server to the other. It is, of course, immediately obvious that such copying ranges from extremely difficult to near-impossible at the volumes of data and velocity of reaction required in this case. A combined operational/informational environment, where ACID transactions can be run simultaneously in parallel with detailed analysis, is clearly required. While many observers associate Google with the introduction of MapReduce and NoSQL technologies, the company has long since moved to their own distributed F1 relational database[8] to perform the processing just described.

This requirement is, of course, not unique to Google. Online retailers face exactly the same challenge: how to decide what to offer prospective customers as cross-sells or upsells based on prior behavior, including the path they have followed during the current session to reach the order confirmation page. All online businesses, from travel sites to banks and insurance brokers, also deal with the problem of combining real-time transactional work with real-time analysis. Their solutions are varied, often depending on the size of the business and the need for speed of reaction, but generally comprise a mix of systems with varying levels of data transfer between them to achieve the balance of ACID characteristics and analytic speed they require.

Integrating data from the Internet of Things

The volume and velocity of events and measures generated by devices connected to the Internet of Things (IoT) are growing exponentially. From smartphones to home appliances, from packaged goods to automobiles, sensors and intelligent controllers are generating a blizzard of data, a large proportion of which demands real-time analysis and instant action. A person carrying a smartphone through a shopping mall provides geolocation data that allows a retailer to know she is approaching their store. The retailer wants to encourage her to buy something, but before sending the offer must analyze in real-time the customer's preferences—privacy status, prior purchases, wants and needs based on social media posts, and more. This analysis of current and historical data, the creation of the offer, and the reaction to it must all be performed in the few seconds or less that it takes the prospective customer to pass the storefront.

Copyright © 2014, 9sight Consulting, all rights reserved

Logistics and transportation companies battle daily to compensate for a variety of real-time challenges—weather, traffic, mechanical failures, etc.—as they match cost containment, service commitments and changing customer needs. The IoT offers a plethora of new, real-time data points that allow business trade-offs; but only in combined operational/informational systems capable of ingesting and curating this flood of semi-structured data and joining and analyzing it with traditional business data.

Financial services at light speed

Banking today, whether retail or investment, is driven by speed and volume of transactions. And perhaps more than any other business, the reliability and consistency of these transactions are vital. Increasingly, however, these transactions are dependent on real-time analysis and understanding of related data. In retail banking, one aspect of particular importance is the detection and prevention of fraudulent transactions before they complete, based on patterns of behavior spread across space and time. Such analysis requires a combination of historical and real-time data and must occur in a sub-second timeframe. Similarly, in investment banking, trading is increasingly real-time and real-time analytics is required to decide which trades should be undertaken.

In both areas, combined operational/informational systems are one of the ways that the banks can perform the analysis of the required current and historical data. Although some data may be stored and pre-analyzed in a data warehouse or specialized data marts, the predictive models thus generated must be applied in real-time against a combination of recent and immediate transactional data. The fundamental characteristic of such systems is that operational and analytic processes run in parallel on the same data.

The worlds of operations and decision making converge

In a minute there is time for decisions and revisions which a minute will reverse.

T.S. Eliot

As we've seen, the original divergence of the worlds of operations and decision making was driven by both business needs and technology limitations. The business need has now clearly changed: businesses do want to track performance on an ongoing basis—often minute-by-minute—and they no longer recognize significant differences between operational and informational needs in cases such as those just described. But, has technology advanced to support these needs?

In general, faster processors, bigger disks, and massively parallel processing, as well as improved software techniques in columnar storage and compression have driven substantial performance gains in both operational and informational systems. But, there are a number of developments that specifically support convergence of operations and real-time decision making.

1. In-memory databases

The design point of traditional relational database management systems was of the data residing and being continuously maintained on spinning disk. This design was, of course, driven by the limited amount of data that could fit in memory and the need to protect data in the event of system failure. There is a difference of multiple orders of magnitude in speed of access to data on disk vs. in memory, leading to a significant bottleneck while shuffling data in and out of memory for processing. Compensating for this bottleneck has long been a central design consideration for RDBMSs.

Copyright © 2014, 9sight Consulting, all rights reserved

The obvious solution is in-memory databases, the use of which has substantially increased over the past few years, driven by the convergence of two hardware developments. First, 64-bit addressing offers the possibility of huge data arrays in memory, in principle, up to 16 exabytes. In practice, current physical and virtual memory constraints may limit this to a multi-terabyte range. However, this enables significant database sizes that can support all but the largest application needs. Second, with memory prices dropping steeply, memory sizes of hundreds of gigabytes to a few terabytes are becoming economically feasible, although still significantly dearer than disk. Much modern database design work thus aims to manage much or all of the database in memory, with longer-term resilience and historical storage supported by disk.

While re-architecting RDBMSs and, in some cases, redesigning applications is a slow and expensive process, the advantages are significant. With such systems, the typical performance trade-offs between an operational, row-optimized, read/write design and an informational, bulk-access, read-only design are substantially reduced. A single data layout can better satisfy both needs. Typical BI performance designs, such as multiple indexes, denormalization and column-orientation, which are expensive to build and maintain, can be minimized.

2. Reducing contention for data access

Traditional database designs used a system of locking records to ensure consistency of the data in cases where multiple transactions try to read or write the same data element. The first transaction simply locks the data element, preserving its state for the duration of its use. In operational systems, locks tend to be short-lived as data is read and updated quickly; in informational systems, on the other hand, queries can be long running, leading to read locks being held over longer periods. These read locks can significantly interfere with operational activities in the database. This behavior was one of the main technical reasons why operational and informational systems were separated, offloading data via ETL or replication tools to a data warehouse, resulting in delayed analytics.

Three fundamentals of combined operational/ informational systems:

- *In-memory*
- *Low contention*
- *Scale out*

An alternative and more modern approach to maintaining consistency is known as multi-version concurrency control (MVCC). Here, data is never physically overwritten: rather, data changes are always appended, allowing multiple versions of the data to coexist in the database. Long-running analytical read-only queries no longer block subsequent write (or read) transactions because they operate only on existing versions of data that will never change. MVCC is now available in most commercial RDBMSs in addition to traditional locking methods.

3. Scaling out the relational database

Traditional RDBMSs, designed for largely operational needs with strong ACID requirements, were first built in an era when computing was performed on centralized, mostly single-processor systems. These databases were designed with one process managing all transaction access to data. Although increasingly complemented with massively parallel processing (MPP), especially in informational systems, the underlying database design remains highly centralized today. With increasing volumes of both data and transactions, such centralization becomes a significant bottleneck, as growth is supported mainly through scale-up to larger and more powerful servers.

Scale-out—the ability to expand a system by adding peer-equivalent nodes—has long been a feature of application servers in the Web ecosystem. However, scaling out databases has proven more difficult, with most approaches requiring an eventual-consistency model (often called NoSQL) that is incompatible with operational and near real-time informational processing. True scale-out of

Copyright © 2014, 9sight Consulting, all rights reserved

RDBMSs requires re-architecting the underlying processing model of the database engine to comprise object-oriented processes that operate independently and collaborate as peers.

As we shall see later, NuoDB addresses all of these points, enabling convergence of the operational and informational environments.

A NOTE ON HYBRID TRANSACTION/ANALYTICAL PROCESSING (HTAP)

In January 2014, Gartner published a research note *"Hybrid Transaction/Analytical Processing Will Foster Opportunities for Dramatic Business Innovation"*, addressing exactly this convergence of operational and informational needs. The acronym, HTAP, echoes the old OLTP and OLAP (On-Line Transaction and Analytical Processing) terms widely used in the '80s and '90s.

On the business side, the report emphasizes the importance of real-time advanced analytics and points to the emergence of real-time business-driven decision-making processes and "intelligent business operations". In technology terms, it focuses on in-memory computing as the key enabler for these developments. The report also provides an impact analysis of these changes and recommendations on how businesses should prepare for them.

REPORTS OF THE DEATH OF DATA WAREHOUSING ARE EXAGGERATED

The convergence of operational and informational systems is essential for high performance, real-time operations that rely on both analytics and transaction processing. However, not all processes have such demanding needs. As a result, converged operational/informational systems do not replace all data warehousing or data mart implementations.

Where there is a need for consistent, reconciled reporting, for example to regulatory authorities, a data warehousing function will continue to provide the highest level of reliability. Such a system offers a single version of the truth and maintains it with historical integrity. However, this new "data warehouse" (labeled *Core Business Information Repository* in figure 2) is likely to be smaller and more limited in scope than today's implementations, which try to satisfy all reporting and consolidation needs. Second, there will continue to exist particular analytical needs which require very specialized data formats to perform quickly or efficiently. As shown in figure 2, these new data marts or *Core Analysis and Reporting Stores* are more likely to be fed directly from the converged operational/informational store rather than from the data warehouse as is recommended today.

Figure 2: A future data architecture

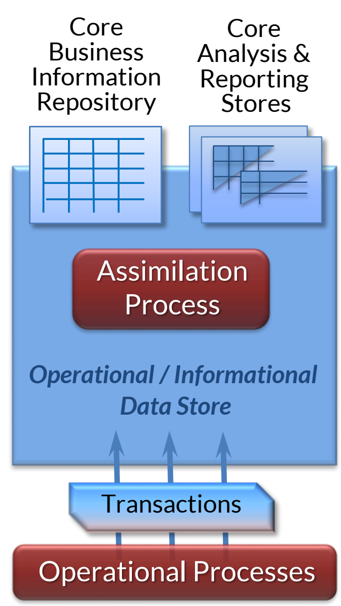

In addition, it is important to note that not all data will reside in the converged operational/informational data store. Not all data requires the level of simultaneous update and analysis offered there. Even in the longer term, some applications may never be migrated into the converged operational/informational environment. Reasons may include the timing, cost or complexity of migration or the lack of a business need for real-time analytics. Integration of data from these systems and the converged operational/informational store will have to occur in the data warehouse, as it does today. This points to a continuing need for assimilation processes to manage data within this environment. Such functions, while likely to be less complex and extensive than today's extract-transform-load (ETL) systems, should be included in the overall design.

These considerations point to a long period of coexistence between current layered architectures and the combined operational/informational environment. New applications, as well as those where real-time analytics are vital, will likely be the initial targets for implementation or migration. Together with the technological features described above,

Copyright © 2014, 9sight Consulting, all rights reserved

this suggests that investigating a new platform that supports the specific needs of combined operational/informational use and applying it in areas of high value real-time use would be an appropriate first step.

What NuoDB offers

Built from the ground up over the past few years, with modern hardware functionality and software insights firmly in mind, NuoDB is a distributed memory-centric database that utilizes a three-tiered architecture: an administrative tier, a transactional tier, and a storage tier. Each tier can be composed of processes running on multiple hosts in one or more datacenters or regions.

NuoDB's transactional tier is composed of one or more transaction engines (TEs) operating on one or more hosts. TEs perform the SQL operations typical of a relational database management system. Applications connect to TEs to access data. TEs aim to keep as much as much of the data being used by the application cached in memory. If an application requests data from a TE that is not already in cache, that TE will request that data from a nearby TE, if the data is cached there, or from the storage tier. Storage tier components, called Storage Managers (SMs), are responsible for providing persistent data storage, and for supplying data to TEs as requested. In ACID database terms, TEs are responsible for Atomicity, Consistency and Isolation, while SMs implement Durability. The administrative tier is a collection of processes that manage the lifecycle of TEs and SMs, and provide the connection points for applications. The processes in each tier can be distributed across multiple hosts in a NuoDB domain. A domain can support multiple databases, and each database can be configured with as many TEs and SMs as needed.

Returning to the three hardware/software characteristics required for combined operational/informational processing outlined in the previous section, let's check how NuoDB stacks up.

NuoDB is a modern, distributed, memory-centric database, composed of processes scaling out over multiple hosts in one or more datacenters, which can address operational/ informational needs.

1. *In-memory database:* NuoDB operates on data with a primary expectation that the data it needs is in memory, initially with the local TE and then with another TE. Only if both these options fail does it go to an SM. The approach is based on a built-in, on-demand caching scheme. All transactional operations are done on in-memory data, giving good performance and reliability guarantees, with durability and data availability provided by SM stores.

   Under the covers, NuoDB is very different from traditional relational databases. It is built on an object-oriented paradigm, with data stored in memory as independent, relatively small "atoms" that are structured to enhance distributed, in-memory operation. SQL operations are provided by an overlying layer: applications speak SQL and need to know nothing of how the in-memory aspects work. Similarly, the SMs isolate in-memory from storage concerns: data can be hardened to disk in any appropriate format or file system.

2. *Reduced contention for data access:* NuoDB uses multi-version concurrency control to manage contention for access to data. This means that long-lived analytical queries on recent transactional data won't block ongoing operational workloads. All workloads see MVCC-mediated access to the latest data at the level of isolation appropriate to the needs of the application or even individual queries.

Copyright © 2014, 9sight Consulting, all rights reserved

3. *Scale-out:* NuoDB is designed to dynamically scale out horizontally to accommodate larger or additional workloads. This means that even as TE nodes are added to increase operational throughput, other TE nodes can also be added in sole support of analytical workloads. NuoDB supports this workload-specific affinity via an easy-to-configure load balancer. As a result, analytic workloads can be directly bonded to specific TEs on specific hardware configurations (such as high-memory for TE hosts). This means that cache pollution is avoided: the in-memory caches for those analytics TEs remain "hot" for analytic workloads, while the caches for operational TEs remain "hot" for those workloads. This is shown in figure 3, where transaction engines are logically grouped to separately address operational or informational needs.

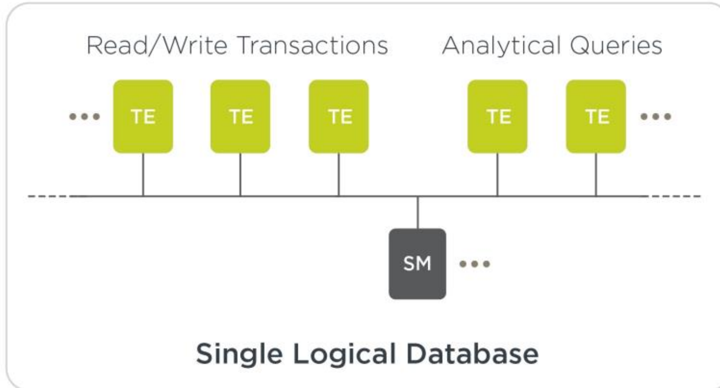

Figure 3: NuoDB operational/ analytical configuration

The outcome is a database that can be distributed over as many servers as needed to provide the performance required. The servers can also be provisioned dynamically, because each TE has access to its own in-memory atoms of data, to those of another TE or, failing that, to a disk-based copy through a storage manager. Operational and informational work can peacefully coexist on the same database with minimal contention between read and write processes. Under the covers, the database takes care of ACID requirements in a fully distributed environment, with the applications using only standard SQL, and completely isolated from the physical processing, memory and storage required to make this all happen. Designed to run on commodity hardware/software platforms, NuoDB can be implemented on premises, in the cloud or in a hybrid environment.

Conclusions

It is only in our decisions that we are important.

Jean-Paul Sartre

Today's business needs for near real-time analysis and action taking demand a new way of addressing operational and informational processing. Gone are the days when copying all data from operational systems into a data warehouse or mart for informational use met all needs. The speed of commerce today, the Internet of Things, and new business models, common across all industries and business sizes, require the use of a combined operational/informational environment. Only such an environment allows real-time analytics to be performed on current and historical data, with the analysis results used to drive immediate and highly specific actions, while traditional operational transactions continue to be processed against the same data. It should be noted that such combined operational/informational systems are likely to complement more traditional data warehousing approaches for the foreseeable future.

Modern business needs for integrated operations and analytics demand a combined operational/ informational system in key business areas.

Traditional RDBMS implementations, with their roots in '70s and '80s technology are often poorly suited to satisfy combined operational/informational needs. Three major characteristics can be identified for such support: (i) the use of a primarily in-memory model to provide sufficient performance, (ii) a flexible

Copyright © 2014, 9sight Consulting, all rights reserved

approach to reducing contention between reads and writes, and (iii) the ability to easily scale out to handle the large volumes of detailed data involved. A new technological base is required.

One such new base technology is provided by NuoDB, a distributed, memory-centric relational database that uses multi-version concurrency control to reduce read/write contention and a novel object-oriented processing model that decouples the SQL model from the underlying in-memory and storage mechanisms, allowing fully flexible scale-out. This approach enables combined operational/informational systems based on standard SQL to be deployed on commodity hardware on premises, in the cloud or in a hybrid model at arbitrary scale, limited only by memory size and processor performance. On this basis, NuoDB offers an ideal environment to begin the journey into the emergent operational/informational world.

NuoDB offers an ideal environment to begin the journey into the emergent operational/ informational world.

*Dr. Barry Devlin is among the foremost authorities on business insight and one of the founders of data warehousing, having published the first architectural paper on the topic in 1988. With over 30 years of IT experience, including 20 years with IBM as a Distinguished Engineer, he is a widely respected analyst, consultant, lecturer and author of the seminal book, "Data Warehouse—from Architecture to Implementation" and numerous White Papers. His new book, **"Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data"** (http://bit.ly/BunI-Technics) was published in October 2013.*

Barry is founder and principal of 9sight Consulting. He specializes in the human, organizational and IT implications of deep business insight solutions that combine operational, informational and collaborative environments. A regular tweeter, @BarryDevlin, and contributor to ITKnowledgeExchange and TDWI, Barry is based in Cape Town, South Africa and operates worldwide.

Brand and product names mentioned in this paper are trademarks or registered trademarks of NuoDB and other companies.

[1] *"What a Wonderful World"*, Bob Thiele and George David Weiss (1967), sung by Louis Armstrong

[2] Devlin, Barry, *"Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data"*, (2013), Technics Publications LLC, NJ, http://bit.ly/BunI_Book

[3] Gartner Press Release, *"Hybrid Transaction/Analytical Processing Will Foster Opportunities for Dramatic Business Innovation"*, G00259033, 28 January 2014, https://www.gartner.com/doc/2657815/hybrid-transactionanalytical-processing-foster-opportunities

[4] *"San Francisco (Be Sure to Wear Flowers in Your Hair)"* John Phillips (1967), sung by Scott McKenzie

5 Power, D.J. *A Brief History of Decision Support Systems*, v 4.0, (2007), http://bit.ly/1IZxHF

[6] *"Dallas"*, an American TV soap opera (1978-91) http://en.wikipedia.org/wiki/Dallas_%281978_TV_series%29

7 Devlin, B. A. and Murphy, P. T., *"An architecture for a business and information system"*, IBM Systems Journal, Volume 27, Number 1, Page 60 (1988) http://bit.ly/EBIS1988

[8] Shute, Jeff, et al, *"F1: A Distributed SQL Database That Scales"*, VLDB (2013), http://research.google.com/pubs/pub41344.html

Copyright © 2014, 9sight Consulting, all rights reserved