And now… a Production Analytic Platform

OCTOBER 2017

A ThoughtPoint by
Dr. Barry Devlin, 9sight Consulting
barry@9sight.com

The data warehouse can and should evolve into a Production Analytic Platform to support operational implementation of predictive analytic models developed in the data lake and elsewhere. The Teradata Database embeds analytic function that supports this goal.

Business and the data warehouse evolve

In an era when business needs access to all imaginable data and the ability to take near-instantaneous decisions and action based on it, the traditional data warehouse database is evolving to provide a foundational platform.

The traditional data warehouse architecture is based on five assumptions[1] that separate operational and informational processing, which suited the data types and business needs of the day. Modern data characteristics (size, structure and speed) and business needs for predictive analytics and immediate action have led to claims of the imminent death of the data warehouse. The future, according to this theory, lies with Hadoop, NoSQL, and more, implemented as a data lake.

The flaw in this thinking is that it forgets the technology underpinning the data warehouse architecture. This technology is a relational database management system (RDBMS) and has been growing in power and evolving in capability for over four decades.

On its own, RDBMS technology has proven its ability to handle the wildly diverse requirements of (1) operational systems, (2) enterprise data warehousing, (3) analytical data marts, (4) personal computers, and (5) open source environments as shown in the table overleaf. In addition, support for newer data types and technologies, such as XML, JSON, and advanced analytic functions such as pathing, sessionization, text analytics, and scoring functions, have been incorporated in RDBMSs, taking advantage of their existing scalability, reliability and performance characteristics.

This pattern of extending existing RDBMSs with complementary function suggests the evolution a new platform, capable of bridging today's hard operational-informational divide. We call this new approach a *Production Analytic Platform*.

SPONSORED BY

www.teradata.com

Requirement	Characteristics	Examples
Operational Systems	*Run the business:* high reliability, up-to-the-second read/write, ACID (Atomicity, Consistency, Isolation, and Durability)	Oracle, IBM DB2
Enterprise Data Warehousing	*Manage the business:* high reliability updating to build consistent history, read-only in daily use	Teradata Database
Analytical Data Marts	*Understand the business:* high speed read-only sorting, summarizing, querying, and reporting	Greenplum, Netezza, Vertica (now acquired by major S/W vendors)
Personal Computing	*Understand the business:* small scale read-only analytical data mart function for experimentation	Microsoft Access
Open Source	Extending all above strategies to the Hadoop world	Hive, HBase, Impala, Kudu, Parquet

Understanding the operational-informational divide

Operational systems record and manage the parties, agreements and transactions that constitute the legal basis of business and must be fully current, consistent, reliable, performant, etc. Informational systems—data warehouses, marts, etc.—contain data extracted and derived only from operational systems to enable trusted analysis, querying, reporting, and decision-making, as shown on the left of Figure 1.

This convenient (for IT) division of labor has been disrupted by two factors. First, human-sourced information and machine-generated data[2] from customers' and partners' digital activities (from mainly external sources) is poorly suited to both traditional operational systems or data warehouses, because of quality, volume, velocity, and variety concerns. Second, large-scale predictive and operational analytics, based primarily on these new sources, demands functionality and scale that characterize the data lake, an environment that is informational in nature and where cheap data storage takes precedence over high performance analytics and predictable user SLAs.

Data scientists first cleanse raw data—an operational task—normalize and analyze it—an informational activity—in the data lake, then pass the data and derived models to the operational environment for ongoing production. This complex flow challenges the traditional operational-informational architecture. In addition, normalization and analysis of this new data often requires the use of reference and other data stored in the data warehouse and operational environments, shown by the gray dotted arrows in Figure 1. The current approach of building a stand-alone, data lake environment for digital activity data and analytics on a different, separate platform to that of traditional business computing thus also creates barriers to such data sharing and reuse.

Figure 1: The operational and informational worlds collide

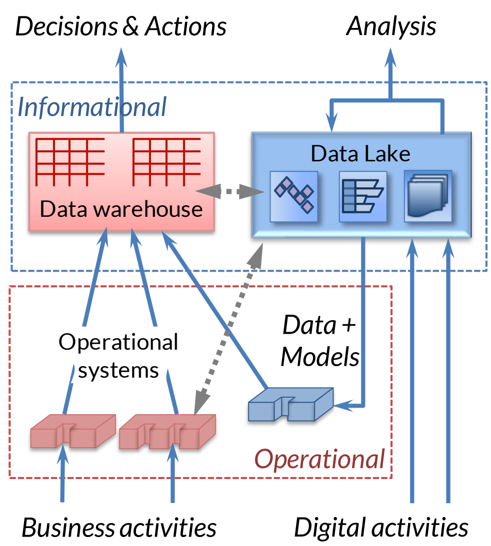

Some data lake proponents favor moving everything into the lake. For enterprises with significant investment in traditional environments, this approach is neither feasible nor financially attractive. Furthermore, the maturity of data lake tools lags behind that of traditional systems, especially in terms of the reliability, availability and maintainability needs of mission-critical operational systems and the data warehouse.

Introducing the Production Analytic Platform

A Production Analytic Platform bridges the traditional operational and informational worlds. Its initial focus is to ease the task of putting models developed in a data lake into a high-performance, reliable environment, but can extend to a range of other borderline operational/informational activities.

As we've seen, the traditional operational-informational divide impedes modern analytics based on digital activity data. Bridging these two environments, the Production Analytic Platform provides a better balance of function and performance between them, based on the eight key characteristics listed here.

Key Characteristics

1. Embedded support for a wide range of reporting, query, and advanced analytics functions, as well as openness to inclusion of new, emerging features
2. Built-in storage and support of all common data formats and the ability to easily include user-defined formats with adequate processing performance
3. Scalability to data volumes necessary for production use of digital activity data (may be less than for initial analytic requirements and must be defined)
4. Ability to access data stored remotely (data virtualization) and to optimize use of such data by local caching and other means
5. Users' ability to access all data types via native languages as well as via SQL
6. Access to multiple analytics engines that provide a choice of tools and analytic methods to users, including commercial and open-source products
7. Support for a wide variety of user types and their preferred tools: business analysts, data scientists, application developers, executives, and business knowledge workers
8. Reliability, availability, maintainability, and performance levels compatible with production use for daily operational decisions, as well as tactical and strategic decision making

The first and central practical question about this platform is to identify its base technology. The most likely candidates are a classic RDBMS and open source Hadoop-related technology. The comparison is between a mature technology that has been designed and built in a controlled and managed approach over decades (RDBMS) and one that is still evolving and that is driven by the immediate and often short-term interests of a development community over less than a decade. The latter traits certainly contribute to rapid technological advances and, for characteristics (1) - (4), means that even where RDBMS is currently ahead, we may assume that open source will catch up and perhaps overtake in a few years. For characteristics (5) - (7), both approaches play fairly evenly.

Characteristic (8), therefore, becomes key. The open source environment presents extensive, well-known software maintenance issues. Rapid feature development often comes at the expense of reliability and operability. In such areas, RDBMS has a considerable advantage, a long history, and is already embedded in the operational and informational environments of most enterprises. And, as discussed above, RDBMSs continue to add support for a variety of novel data formats and analytical functions.

The Production Analytic Platform is thus positioned across the operational-informational boundary, as shown in Figure 2. It allows direct ingestion of digital activity data to support production analytics based on models generated previously in the data lake. With a shared store of core business data from the traditional operational systems and the data warehouse, the Production Analytic Platform is the ideal location for ancillary data and function used in analytics both in production and in the model-build phase in the data lake.

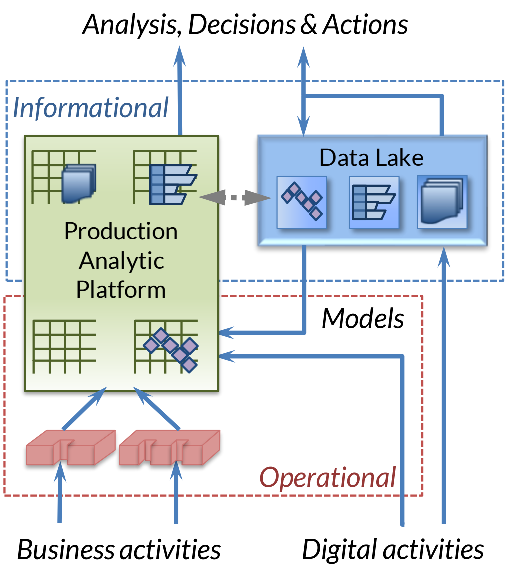

Figure 2: Production Analytic Platform

In terms of implementation, the Production Analytic Platform is an extension of the data warehouse (and, in particular, the enterprise data warehouse) to undertake more operational activities. These activities are firstly related to day-to-day operation of predictive analytic models, but can extend to other borderline operational/informational activities. Central to all of this is the analytical feature richness, robustness, and performance characteristics of the RDBMS underpinning the data warehouse, all of which points to the Teradata Database as a suitable starting point.

Further ThoughtPoints in this series explore specific aspects of the analytic extensions in the Teradata Database that support this direction.

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 book, ***"Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data"*** *(http://bit.ly/BunI-TP2) was published in October 2013.*

Barry is founder and principal of 9sight Consulting. He specializes in the human, organizational and technological implications of deep business insight solutions combining all aspects of internally and externally sourced information, analytics, and artificial intelligence. A regular contributor to Twitter (@BarryDevlin), TDWI, BACollaborative, and more, Barry is based in Cape Town, South Africa and operates worldwide.

[1] Devlin, B., *"Business Integrated Insight BI2—Reinventing enterprise information management"*, (2009), http://bit.ly/2xZc77b

[2] Devlin, B., *"Business unIntelligence—Insight and Innovation beyond Analytics and Big Data"*, (2013), Technics Publications LLC, NJ, http://bit.ly/BunI-TP2

Production Analytic Platform—It's a Matter of Time

NOVEMBER 2017

A ThoughtPoint by
Dr. Barry Devlin, 9sight Consulting
barry@9sight.com

The Production Analytic Platform marries predictive analytics to production values and goals. Handling the complexities of analyzing time dependent data, especially from the Internet of Things, is central to this objective. The Teradata Database embeds extensive temporal and time-series storage and analytic function to enable success.

Analytics of IoT data is all about time

The analysis, representation, storage, and management of data as a function of time is vital as the Internet of Things (IoT) delivers ever more data. The Production Analytic Platform plays a key role in providing this functionality.

Time waits for no man, according to an old proverb. In the modern world, the Internet of Things waits for nobody. In the early days of data warehousing, an end-of-week snapshot was good enough for management; overnight was state of the art. Today, many decisions driven by the IoT and other real-time sources are worthless if delayed at all. Furthermore, such data often makes sense only if analyzed in *time series* that provide the context of what preceded and follows it. Traditional data types and decisions based on historical periods of time—weekly, monthly, etc.—also continue to be important.

These trends—increased data timeliness in a growing set of decisions as well as a broader range of timescales of data for analysis—amplify the demands on technology to handle time fully and efficiently. From database storage to application design, the issue of time becomes central.

Timeliness and other temporal considerations were important in delineating operational and informational environments. However, in the current business and technical world, such a simple binary distinction does not suffice. The Production Analytic Platform is designed to offer a broader range of possibilities. Two key aspects are explored here: (1) temporal database structure, which specifies how time is represented and manipulated in the data stores, and is a foundation for (2) time series support, which details how changes in data over time can be meaningfully captured, managed and analyzed.

SPONSORED BY

www.teradata.com

Prime time for analytics

A Production Analytic Platform provides the ideal locus for time series analytics because of its combination of a wide range of powerful analytic function and its high performance and highly reliable operational characteristics.

If you were a banker, you'd see sums of money flowing in and out of your accounts. Each movement is a *transaction*: an event of legal significance that happens at a moment in time. Although you are quite happy to track all these transactions, managing the business requires, in addition, a different view, such as the total balance in your accounts as of any given time or period. This *status* view, in contrast to a transaction, has a duration—a beginning and end date and time—and is central to business computing in every industry.

The data generated by Internet of Things records events, measures and messages[1] from the real world that are similar in structure to financial transactions. All these types of data are examples of time series data. Our focus here is on IoT data, because its size and speed give rise to particular challenges as well as offering new business opportunities.

Time series become serious business

A Boeing 777 starts up its engines: two *events* are recorded a few minutes apart. A stream of events follows: ailerons set, brakes released, and so on. The captain speaks to the control tower: a *message*. Meanwhile, on-board computers are monitoring engine temperatures, fuel flows, voltages and more on a sub-second basis: *measures* by the million. From now until the flight ends, time series are unfolding measure by measure, event by event, message by message, creating a comprehensive record of the flight.

This data, terabytes per flight, offer airlines and aircraft manufacturers, airports and aviation authorities enormous business opportunities. Preventative maintenance avoids costly flight cancellations with attendant airport disruption. New analyses reduce fuel consumption. Safety standards are improved. Similar opportunities for cost saving and profit generation are evident in almost every industry. The Internet of Things is revered; actually, it's the time series data it produces that is the real boon.

Time series data is nothing more than a time-ordered set of data records, each of which carries the date and time (usually) of creation. While common in scientific computing, times series data has been relatively rare in business computing until the advent of the IoT. Each record consists of a timestamp, an identifier of its recording device, and a payload of data of interest, often little more than strings of numbers or text. In the case of measures, observations are made at regular intervals; events and messages are recorded when they happen. Even with regular observations, data records may be lost in transmission. As a result, the vast majority of IoT data is classified as irregular time series.

Each sensor creates its own time series of identically coded measurements. A thermal sensor provides a single temperature in each record, for example. Controllers aggregate data from multiple different types of sensors into a single record, meaning that the payload may consist of a string of values of different types, such as temperature, velocity, location, etc., called multivariate time series. Because the content can change over time, the payload is often stored in flexible formats, such as JSON, CSV, and so on.

Analysis of time series data begins with the basic observation of trends. Is the temperature of this engine part increasing? When might it reach a critical value? On its own, this may be interesting, but real analytics only begins when multiple time series can be compared. Is there a simultaneous rise in vibration in an adjacent component? Can it be correlated with a change in speed or rapid ascent? What about weather conditions?

Due to record timing problems, missing values, and so on, most time series analytics is based on grouping records into time buckets or intervals, the starting point and size of which is determined by differing analytical needs.

It is the management and manipulation of multiple time series in flexible bucketing schemes that challenges data scientists. Within a single time series, simpler analytical functions—such as the mean of a sliding window—can be computed manually, paying careful attention to missing values and other issues. However, when data consists of multiple time series, each with different time intervals and missing values, the challenges rapidly escalate. Specialized, time-series aware, in-built function is the only viable approach.

Flat file solutions run rapidly out of control and power as data volumes increase. NoSQL data stores are a better solution in terms of scalability and performance, but often suffer from their weak metadata management and limited attention to operational reliability and maintainability needs.

The Production Analytic Platform to the rescue

The Production Analytic Platform offers the best of all worlds by combining the power and management of a relational database with the ability to store non-relational data, and manipulate and analyze it with a combination SQL and advanced analytic functions.

The Teradata Database now offers this foundational function with the recent introduction of support for time series data. This includes the ability to load and store data in specialized tables with a primary time index and a set of time-aware SQL functions operating within and across time buckets. Existing function that supports non-relational data, such as JSON, with full SQL support allows for multivariate time series.

In most instances, time series analytics is a high-skill exercise, requiring data scientists with knowledge of specialized languages and/or techniques. By embedding the function in the Teradata Database, access to the function is simplified and broadened. Straightforward SQL statements allow a broad range of "ordinary" business people to easily modify analytic parameters or set and change the size of the time buckets. Analyses can be iterated easily and quickly, leading to faster decisions and more relevant actions.

The space-time context

For business, two aspects of time are of interest. Time series data from the IoT captures the dynamics of the changing world. Status data from the operational environment provides the context in business terms.

But there's more! At the top of the last section, I mentioned that the second type of time-aware data, status data, was central to business computing. With the hype around the IoT and the focus on time series data, we must keep this traditional data in mind because it is the primary source of context for time series data in the business.

Consider again our Boeing 777 airplane. Analysis of its flight data predicts that an auxiliary power unit is likely to fail within the next 50 hours of flight. It's easy to see that it should be replaced before it fails, but it's also not a safety issue, so the decision for the airline is when and where should the maintenance be performed. This depends on understanding the plane's schedule, the location and availability of the part, the cost of using a standby aircraft if necessary, and a host of other business considerations.

The data required for this decision process resides in the traditional operational systems and data warehouse of the airline—all status data with an important temporal aspect: the time periods during which all the above data is valid. The underlying technology is known as a temporal database and is a key component of the Production Analytic Platform.

Temporal databases (as the) rock

Data modelling and database design start from a view of the world in some mythical moment of "now" and ensure that the current business status is fully described and its interrelationships accurately recorded. In the real world, of course, to this is added business transactions that change the status. Furthermore, mistaken entries are made and must be corrected. Technical glitches corrupt data and are rectified later. All these events result in complex application design and extensions to the data model and database design. Recording time for events and transactions is simple. Status data is more challenging.

Tom Johnston provides an in-depth explanation of the true philosophical and technical complexity of this status data[2]. Since the early '90s database designers have proposed that time in status data is best represented in a ***bitemporal*** model, which adds two timestamps to each database record: ***valid time***, during which a fact is true in reality and ***transaction time***, during which the database record is accepted as correct. Indeed, Johnston argues that further timestamps may be needed in certain circumstances. However, even bitemporal support has been slow in its implementation, arriving only in the current decade, and limited to a subset of the main databases.

With its focus on data warehousing, the Teradata Database became the first mainstream database to add support for bitemporal data in 2010 and has since been upgraded with enhanced temporal analytic features such as derived periods and sequenced views. The correct handling the temporal aspects of status data becomes even more important in the Production Analytic Platform because the enormous volumes and variety of events to be processed demands a solid foundation of status data and its temporal context.

Conclusion

A key consideration for the Production Analytic Platform is its support for both time series data and the time-enabled status data from the operational environment and data warehouse that provides the business context.

Time series data is the basis for advanced analytics of the output of the Internet of Things. Bitemporal status data describes the fundamentals of the business, from logistics to financials, and thus offers the context in which business decisions and action can be taken based on time series analytics. The ability to fully support both these aspects of time—time series data and bitemporal data—in a single environment is a central requirement in the design and implementation of a Production Analytic Platform.

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 book, **"Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data"** *(http://bit.ly/BunI-TP2) was published in October 2013.*

Barry is founder and principal of 9sight Consulting. He specializes in the human, organizational and technological implications of deep business insight solutions combining all aspects of internally and externally sourced information, analytics, and artificial intelligence. A regular contributor to Twitter (@BarryDevlin), TDWI, BACollaborative, and more, Barry is based in Cape Town, South Africa and operates worldwide.

[1] Devlin, B., *"Business unIntelligence"*, (2013), Technics Publications, New Jersey, http://bit.ly/BunI-TP2

[2] Johnston, T., *"Bitemporal Data: Theory and Practice"*, (2014), Morgan Kaufmann, Waltham, MA, http://bit.ly/2gmpdDR

Production Analytic Platform—A Shrinking Decision Cycle

December 2017

A ThoughtPoint by
Dr. Barry Devlin, 9sight Consulting
barry@9sight.com

A blend of predictive analytics with production values and goals is a mandatory foundation for the tightly-coupled, closed-loop decision-making cycles characteristic of modern digital businesses. As time-to-decision decreases, reliability, maintainability, and other qualities of the Production Analytic Platform become increasingly important.

From discovery to action—and back

The MEDA model defines a closed-loop cycle of decision making and action taking—function that has been implemented in a siloed manner in today's analytic world. The Production Analytic Platform breaks down these silos and improves analytics across the whole business.

It's many years, perhaps decades, since it has been acceptable to suggest that a decision maker should go get a cup of coffee while the required data is being crunched. Of course, some decisions do demand considerable time for thought and some data is so big that crunching it is far from instantaneous. However, modern business must operate in cross-functional, tight and closed loops, avoiding any potential delay that introduces the possibility of being outmaneuvered by a competitor or ditched by a customer.

Nonetheless, much design thinking still starts from the simplicity of the silo: Optimize for performance of a single, well-scoped function—an app—with maximum control of the needed data and minimized external dependencies. This approach remains much favored by developers because it increases their chances of successful project delivery. However, the longer-term success of the business process is endangered by *ad hoc* data hand-offs, mismatched function, and time-devouring breaks in continuity.

A closed-loop, sense-and-respond approach is required, such as the MEDA model I have long promoted[1]. The acronym stands for:

- *Monitor* what is happening both within and outside the enterprise
- *Evaluate* implications and consequences, and possible actions

- *Decide* among possible/recommended courses of action
- *Act* to change behaviors, and/or processes and link back to monitoring

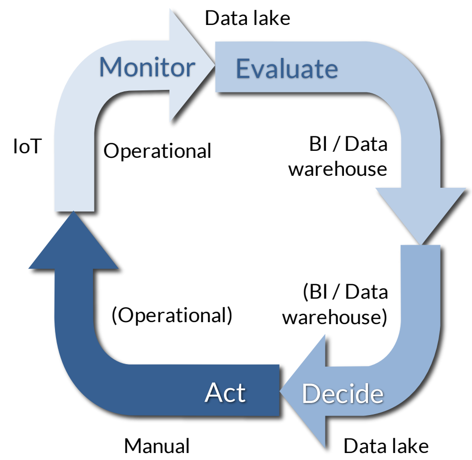

In traditional data warehousing, BI tools support this cycle—mostly the evaluate phase—but much work and coordination is manual. Analytic environments, as typically built today, show more automation but, with the IoT and data lake, lead to a more technically fragmented environment. The figure to the right shows how multiple systems support MEDA.

In the past, monitoring occurred only in operational systems. For production systems, this remains the case. As events in social media and IoT are of increasing interest, they are monitored in the data lake. Depending on the source, evaluation occurs in both the data warehouse and lake. Decision is increasingly automated for faster turn-around, and is also split between warehouse and lake. Action occurs in operational systems where possible or may be manually initiated.

As discussed in the first ThoughtPoint of this series[2], the current approach leads to complex data flows between the three environments—operational systems, data warehouse and data lake. The Production Analytic Platform reduces this data flow complexity by enabling all four phases of the MEDA model within an integrated and performant environment. And, as described in the second part of the series[3], it also addresses the temporal requirements of handling times series and bitemporal data.

With decision making and action taking in MEDA under ever tighter time constraints, two further characteristics of the Production Analytic Platform emerge. First is the principle of bringing the analytic function to where the data resides rather than moving the data to a tool with the right function. The Production Analytic Platform must incorporate a complete set of analytic function, all the way from simple aggregation to advanced analytics, from mathematical & statistical functions to machine learning. Second, the decision cycle must support a wide variety of users with a range of skills and tool preferences. Business analysts review data and make decisions, likely with SQL tools. Data scientists explore data with R/Python, for example. Developers operationalize the results with a combination of tools. The Production Analytic Platform must offer all needed tools.

We now examine the data and analytic function provided by the Production Analytic Platform, based on the MEDA model above.

From data to action on the Production Analytic Platform

A Production Analytic Platform offers multiple data storage formats and a wide range of powerful analytic function in a consolidated relational environment with strong operational characteristics across the full MEDA cycle.

Externally sourced data from the IoT and information from social media arrives in the enterprise in a variety of mostly simple, text-based formats. Internally sourced production data is almost exclusively relational. The Production Analytic Platform must therefore efficiently store this wide array of formats and offer appropriate functionality to process and analyze them with ease. This begins with the Monitor phase of MEDA.

Monitoring the world—within and without

A high-speed train departs from Madrid, its final destination Copenhagen, with only two stops *en route* in Paris and Frankfurt[4]. Soon reaching its top speed of 250 km/hour, it races across the Spanish plains, spewing not smoke but data describing every aspect of its performance. This data—multiple streams of events and measures—arrives continuously at the data processing center of Trans-European Rail (TER) in Berlin, where the progress and performance of all its rolling stock is monitored in real-time.

As discussed previously, this is time series data. TER's specific focus is on the payloads of these streams, containing detailed records of axle rotation speeds, temperatures, vibration levels, and lubrication pressures for every wheel of the train, as well as a range of events on the journey, such as applying brakes, accelerating, and crossing railroad switches. With multiple sensor and edge processor manufacturers, such data arrives in many formats, from simple comma separated variables and key-value pairs to sophisticated JSON and AVRO structures, in text and binary modes. To effectively monitor performance, the Production Analytic Platform stores all these formats natively within the relational database with equal ease. Using SQL with appropriate extensions, business users and applications monitor the incoming data streams and flag any unexpected events or measures. And with more traditional operational data in the same environment, any impacts on passengers or schedules can be easily spotted.

Evaluation and decision—models, analysis, and answers

In an environment where externally sourced data—be it from social media or the IoT—is the basis for evaluation and decision, only statistics and model-based analytics can operate at the scale and speed required. In the case of TER, the data exhaust from their high-speed train offers the opportunity to observe in near real-time the performance of their equipment, to note emerging problems, and to predict—and avert—potential failures.

The starting point is to develop models that correlate observed abnormal measures, such as increasing temperatures or vibration levels in the absence of braking events, with known failures in axle assemblies and subsequent train breakdowns. Such models, usually produced by data scientists in high-performance analytic data lakes, combine data from multiple time series over many train journeys to predict time to failure in many, varied situations. And one of these situations is emerging on TER's premier Madrid-Copenhagen service that has crossed into France just after midnight. Hydraulic pressure has risen while temperature is beginning to climb on bogie C of carriage 12.

To understand the potential implications of these changes, TER must run an array of models on the near real-time train data that is arriving on its Production Analytic Platform in Berlin. These models originate in the data lake and are maintained and enhanced there. They are executed on the Production Analytic Platform using the incoming IoT data using a variety of analytic tools.

For example, the recently introduced Teradata Analytics Platform, based on a combination of Teradata and Aster technologies, offers R, SAS, Jupyter, and KNIME analysis environments, and will include the Spark and TensorFlow engines in the near future. More tech-savvy analysts can work in Python or SQL as they desire.

But what of the Madrid-Copenhagen train? The news is not good: there is a 52% probability of failure by the time the train reaches Paris, rising to 78% by the time it arrives in Frankfurt. The carriage will have to be pulled from service before it arrives at its destination. But where?

Action—where the wheels meet the rail

Taking action demands data from a much broader set of sources than simply the IoT data from the train. A range of traditional business data will play into determining the best action to take. Occupancy of the affected carriage and availability of equivalent seating elsewhere on the train is an important consideration. In the event of there being insufficient alternative seating, a replacement carriage will need to be attached to the train. The immediate question is where is the nearest one and what knock-on effects on other trains might ensue. Or perhaps a temporary repair might suffice, in which case the locations of components and skilled staff will be required.

A key strength of the Production Analytic Platform is that such traditional operational and informational data is part of the same environment. With such data readily at hand, analyses based on complete information are easily undertaken and can be run repeatedly as updated predictions on likely time to failure are updated from the analytic models of ongoing real-time data from the train.

With hybrid row/column storage, in-memory optimization and vectorization, and automatic multi-temperature data management, as well as an intelligent cost-based query optimizer, the Teradata Analytics Platform provides an ideal environment for such complex and time-sensitive workloads.

Conclusion

> Today's decision-making cycle demands tight integration of very different types of data and function, from collection and analysis of external data, through predicting future states, to taking immediate action based on ongoing operations. A Production Analytic Platform is vital to meet these demands.

Modern business needs for timeliness and cross-organization coordination drive a shrinking decision-making cycle. Real-time data from the external world—social media and IoT—is the foundation for ongoing modelling and analysis of the ever-changing behaviors of people and machines in the physical world. Seamlessly combining such data with traditional operational and informational data from internal systems is vital to ensure a closed-loop MEDA-style decision cycle. A Production Analytic Platform offers an ideal environment to bring the varied data and processing needs together with the required reliability, scalability, maintainability, and performance.

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 book, ***"Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data"*** *(http://bit.ly/BunI-TP2) was published in October 2013.*

Barry is founder and principal of 9sight Consulting. He specializes in the human, organizational and technological implications of deep business insight solutions combining all aspects of internally and externally sourced information, analytics, and artificial intelligence. A regular contributor to Twitter (@BarryDevlin), TDWI, BACollaborative, and more, Barry is based in Cape Town, South Africa and operates worldwide.

[1] Devlin, B., *"Business unIntelligence—Insight and Innovation beyond Analytics and Big Data"*, (2013), Technics Publications LLC, NJ, http://bit.ly/BunI-TP2

[2] Devlin, B., *"And now… a Production Analytic Platform"*, October 2017, http://bit.ly/2k3QmQx

[3] Devlin, B., *"Production Analytic Platform—It's a Matter of Time"*, November 2017, http://bit.ly/2ABTpng

[4] Sadly, in the real word, this journey would require several high-speed trains operated by different companies.

Production Analytic Platform—Into the Future

JANUARY 2018

A ThoughtPoint by
Dr. Barry Devlin, 9sight Consulting
barry@9sight.com

As might be expected of any architectural thinking, the Production Analytic Platform must look beyond current issues and solutions to how business and technology environments may evolve. The final ThoughtPoint in this series considers how advances in analytics and artificial intelligence influence this approach.

A thin line between analytics and AI

The future of analytics is closely aligned to artificial intelligence. In fact, given that both are simply ways of building and exercising models of the real world, the boundary between them is paper thin and growing more porous daily.

Thirty years ago, the first Data Warehouse architecture[1] emerged from the convergence of business demands for improved access to consistent management information and the new possibilities offered by relational technology. With it, BI was born and has become the lens through which decision making is still viewed. Some ten years ago, advances in business needs and technology combined to shift our focus again—to big data, predictive analytics and Hadoop.

Now, even as the foundations for fully operational analytics are being laid with the Production Analytic Platform[2], technology is undergoing yet another tectonic shift: Artificial Intelligence (AI) is supplementing and, in some case, supplanting BI and analytics at the heart of decision-making support.

AI—and an array of partial and overlapping synonyms[3] such as *machine* and *deep learning*, *cognitive computing* and *neural networks*—is a very broad and often confusing topic, ranging from theoretical mathematics, through neurobiology, to fundamental computer science. For anyone with a BI/analytics background, one simple description is as technology that enables machines to learn about and model the real world, to make decisions about it, and act accordingly, minimizing the level of human intervention in the entire process. Autonomous vehicles provide an obvious and easily understood example of both independent learning and automated decision making / action taking.

This explanation emphasizes the thin line between analytics and AI. There are, of course, differences in techniques and tools, for example in the use of neural networks and platforms such as TensorFlow and OpenAI. However, the intent of the AI process is, in fact, an extension of analytics: From searching for the "right" answer among trends to discovering intelligent options and ranked possible solutions to real world challenges. The continuity of this spectrum of function can be seen, for example, in the range of function available in the Teradata Analytics Platform— from SQL to advanced analytic techniques like machine learning, graph, pattern, pathing, and sentiment, to full AI tools such as TensorFlow and Spark on the near-term horizon.

AI and the Production Analytic Platform

As the Production Analytic Platform evolves, AI will become more important, complementing current operational and analytic aspects. Indeed, the evolution of AI will, in part, drive the future of the Production Analytic Platform.

With such a thin line between analytics and AI, it is immediately clear that the Production Analytic Platform must and will include elements of AI libraries and platforms. Even a high-level review of AI tools[4], shows significant overlaps between traditional mining and new deep/machine learning function in many available platforms and libraries.

The emerging Production Analytic Platform therefore spans the entire spectrum of IT processing—from operational systems, through BI and analytics, all the way to AI. It brings together data and function from the initial collection and cleansing of data and information from every environment, internal and external, through simple reporting and problem solving, to deep analytics, both predictive and prescriptive. It this way, the inclusion of AI in the Production Analytic Platform helps close the MEDA loop described in part three of this series[5].

As AI-supported function becomes more common, analytics and basic AI become mere table stakes in business competitiveness. New business needs emerge, driving further advances in technology that enable enterprises to differentiate. This symbiotic relationship between business and IT is the ***biz-tech ecosystem*** described in *Business unIntelligence*[6] that is the foundation of digital transformation.

Future AI also drives the Production Analytic Platform

The significant advances seen in AI in the current decade—particularly in image recognition and natural language processing—rest on two key foundations: (i) the exponential growth in processing power and data storage of modern hardware and (ii) the even faster growth of data/information training sets available from the Internet. Together, these trends have led to significant centralization of AI processing with the problems entailed in the transport of such data from the edges of the network to the center.

Recent developments in AI are focusing on reducing both the training data required and processing power needed. DeepMind's AlphaGo Zero game, for example, learned to play without *any* training data and uses considerably less power than its predecessor[7]. While still in the research lab, the direction for business AI software is clear.

In addition, vendors are increasingly pushing AI function—both training and operation—to the edge of the network. With localized function in smartphones and autonomous vehicles, business users experience enhanced decision-making timeliness and IT can significantly decrease data transportation costs and delays.

It follows, therefore, that these complementary business and IT needs will drive the Production Analytic Platform toward an increasingly decentralized and distributed foundation. This direction is already visible in the Teradata Everywhere™ approach which provides the full power of the Teradata Analytics Platform and open source software across an array of deployment options, from on-premises to the edge of the public cloud, built on Teradata or commodity hardware.

Conclusion

Decision making and action taking today and into the future demand tight integration of very different types of data and function, from collection and analysis, through predicting future states, to taking immediate, operational action. The Production Analytic Platform, based on relational database technology, offers the most appropriate and effective solution.

As a modern logical architecture, the Production Analytic Platform addresses significant changes in data usage patterns and technological advances over the past half-decade. Further change in both business and IT is expected and the platform will evolve to meet this ongoing transformation. The key aspects of this architecture are:

- *Relational database technology* at the core—for reliability, availability, scalability, maintainability, and performance levels spanning operational and analytical usage
- *Built-in non-relational support* for storage of data in a wide range of formats (e.g. JSON, Avro, etc.) and processing approaches from simple query, through analytics and AI, accessible via SQL as well as native languages by users of all types and skill levels using their preferred tools
- *Data virtualization support* for access to and use of data stored remotely in both relational and non-relational formats, optimized by caching and other means
- *Distributed and decentralized operation* to support the longer-term migration of AI to the edge of the network with local data and processing

The four characteristics above are not new. Today, however, they often exist across different products and separate platforms, leaving enterprises to employ or adopt the role of systems integrators. The Production Analytic Platform integrates these functions in a single platform.

Users of analytics and AI are becoming ever more varied in their goals, use more data of different types and sources, and demand faster responses to both operational and informational processes. Indeed, individual users increasingly require seamless transmission of information and expertise from developer to data scientist to business analyst, and *vice versa*. Operationalizing experimental discovery work demands simple and elegant transitions between these roles. Existing IT environments—a mix of operational systems, data warehouses / marts, and data lakes that have grown historically more complex in management and use—are unable to meet these expanding business demands.

Adding yet another technology environment to this mix is not an answer. Rather, a more realistic and achievable approach is to choose an existing environment and expand and improve its functionality to address these increasingly central business needs. The wide-ranging strengths and maturity of the relational database—especially as it has evolved to support enterprise data warehousing—offer the best starting point for this journey. The Production Analytic Platform defines and describes the starting point and the target, both in the immediate future and in the longer term. Better still, the technology already exists to start that journey today.

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 book, **"Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data"** *(http://bit.ly/BunI-TP2) was published in October 2013.*

Barry is founder and principal of 9sight Consulting. He specializes in the human, organizational and technological implications of deep business insight solutions combining all aspects of internally and externally sourced information, analytics, and artificial intelligence. A regular contributor to Twitter (@BarryDevlin), TDWI, BACollaborative, and more, Barry is based in Cape Town, South Africa and operates worldwide.

[1] Devlin, B.A. and Murphy, P.T, *"An architecture for a business and information system"*, February 1988, IBM System Journal, 27(1), http://bit.ly/2hfgiVl

[2] Devlin, B., *"And now… a Production Analytic Platform"*, October 2017, http://bit.ly/2k3QmQx

[3] Vorhies, W. *"Machine Learning—Can We Please Just Agree What This Means"*, December 2017, http://bit.ly/2yh11J3

[4] Davis, S. and Baker, J., "Top 12 AI Tools, Libraries, and Platforms", September 2017, DZone, http://bit.ly/2AwQgbS

[5] Devlin, B., *"Production Analytic Platform: A Shrinking Decision Cycle"*, December 2017, http://bit.ly/2AHSgsE

[6] Devlin, B., *"Business unIntelligence"*, (2013), Technics Publications, New Jersey, http://bit.ly/BunI-TP2

[7] *"AlphaGo Zero: Learning from scratch"*, DeepMind blog, October 2017, http://bit.ly/2zQaEQW