Volume 32 Number 6
DECEMBER 2018/JANUARY 2019

The Journal of
Information Integration and Management

database
TRENDS AND APPLICATIONS

TREND-SETTING PRODUCTS for 2019

WWW.DBTA.COM

Multiplatform Data Quality Management

Unison – Speedy, Secure, Scalable

Melissa's Unison is a data steward's best friend. It's the ideal multiplatform solution to establish and maintain contact data quality at higher speeds – processing 30+ million addresses per hour – while meeting the most stringent security requirements. With Unison, you can design, administer and automate data quality routines that cleanse, validate and enrich even your most sensitive customer information, as data never leaves your organization. Streamline data prep workflows, reduce analytics busy work, gain more insights and increase efficiency!

- Verify and standardize names, addresses, emails and phone numbers, plus append lat/long coordinates and Census data.
- On-Premise platform requiring no coding or programming with automatic updates from Melissa.
- Works offline, scalable across multiple servers and allows users to script batch jobs with various levels of data access.
- Access, manage and visually analyze the quality of your data over time, enjoy project collaboration and more!

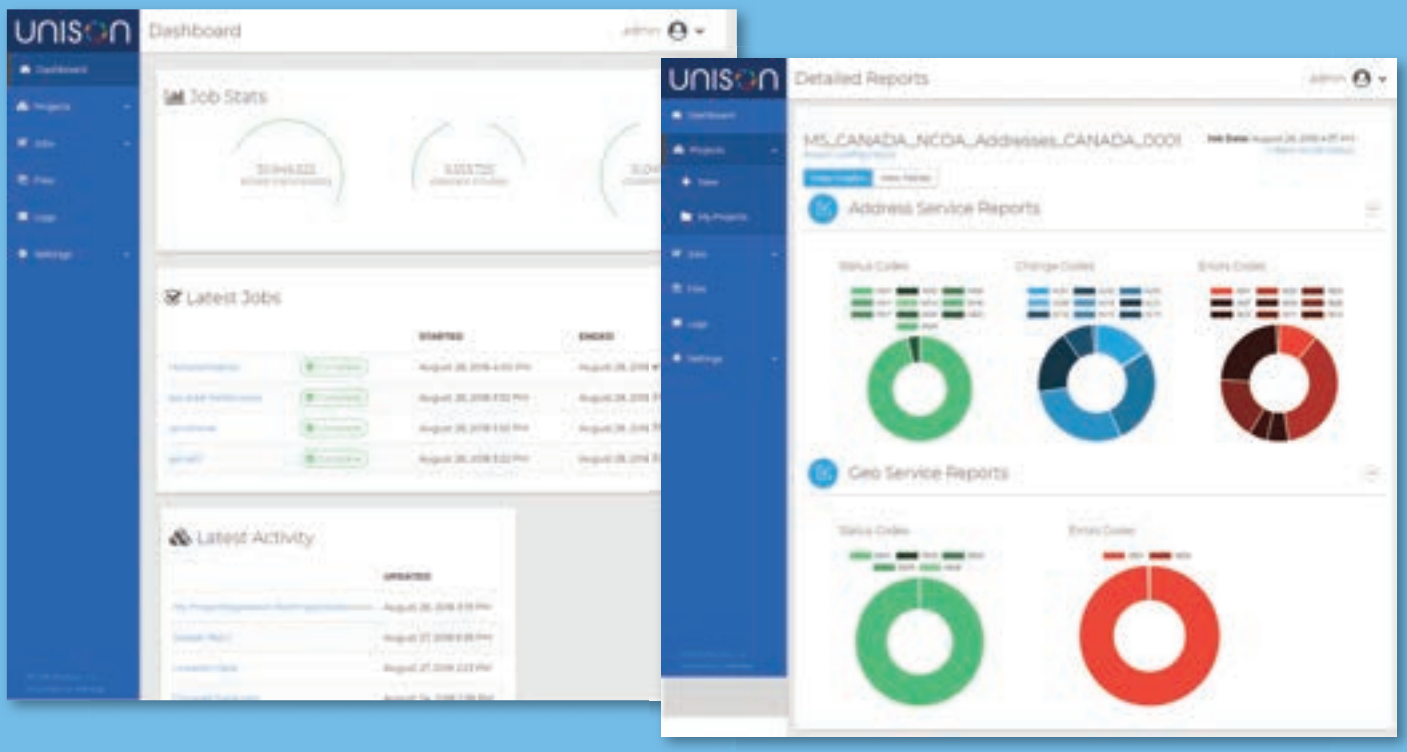

Request a Demo

Melissa.com/dbta-unison

1-800-MELISSA

melissa®

The Journal of Information Integration and Management

database
TRENDS AND APPLICATIONS

Volume 32 | Number 6
DEC 2018/JAN 2019

CONTENTS

PUBLISHED BY Unisphere Media—a Division of Information Today, Inc.
EDITORIAL & SALES OFFICE 121 Chanlon Road, New Providence, NJ 07974
CORPORATE HEADQUARTERS 143 Old Marlton Pike, Medford, NJ 08055

DATABASE TRENDS AND APPLICATIONS (ISSN: 1547-9897; USPS: 16230) is published bimonthly (Feb./Mar., Apr./May, Jun./Jul., Aug./Sep., Oct./Nov., and Dec./Jan.) by Unisphere Media, a division of Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055 USA; Phone (609) 654-6266; Fax (609) 654-4309; Internet: infotoday.com. Registered in U.S. Patent & Trademark Office. Periodicals postage paid at Vincentown, NJ, and additional mailing offices.

FEATURES

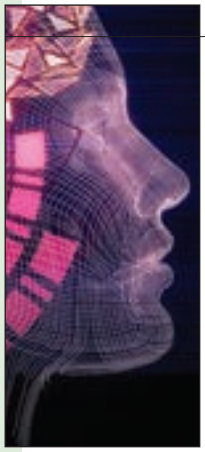

DEPARTMENTS

COLUMNS

MEDIA PARTNER OF THE FOLLOWING USER GROUPS

Cloud Opens the Path to ...

By Joe McKendrick

How fast and far can databases grow, and how can such growth be sustained? That's the question faced by many data managers these days, who deal with growing demands from their businesses for real-time, analytical capabilities, incorporating data-driven initiatives such as the Internet of Things and artificial intelligence. They are responding and keeping up with these requirements through a combination of cloud resources and automation.

These are the findings of a recent survey of 260 data managers, fielded by Unisphere Research, a division of Information Today, Inc., in partnership with VMware and the Independent Oracle Users Group.

The growth of data-driven enterprises is pressuring data administrators to deliver high-performing and responsive systems that can scale with the business. However, many enterprises are encumbered by the licensing and support issues that typically accompany database systems, resulting in potentially high and unexpected costs, as well as skills shortages. While enterprises are turning to the cloud and automation solutions to enhance their capabilities in backup and recovery, the challenge is many data managers subscribing to cloud services are not making licensing costs enough of a priority.

The survey found that database environments are sizable and complex these days. Respondents manage multiple databases, and many have databases scaling into the multi-terabyte stage. A large portion of respondents, 46%, also reported that their largest database exceeds 1TB and range up to 25TB in size.

As organizations keep growing and expanding their data environments, they run into obstacles. The survey found licensing and support is the number-one challenge for organizations seeking to expand the number of Oracle databases and applications. Data may have become the fuel driving today's and tomorrow's enterprises, but regardless of where it comes from or where it is stored—managing it in commercial databases is still tied to the traditional licensing and support model that has been in place for decades. Activating new databases, processor cores, or end user licenses means additional costs. More than four in five respondents reported that it is difficult to grow their data environments due to obstacles with licensing and support from their database vendors—a number that has increased since the initial survey in 2014. The percentage of data managers citing challenges with licensing and support costs, 81%, was up 22 percentage points over the previous survey.

Other issues also stand in the way of growth. Individuals with database skills are getting harder to find. Additional issues are also coming to the fore since the 2016 survey. For example, respondents who cited challenges with finding the right skills has more than doubled since 2016—from 32% to 66% reporting issues. There also has been a surge in administration costs and complexity getting in the way of data environment expansion—cited by 55% of respondents for an increase of 24 percentage points. Of the hardware categories, the only category that showed growth was storage cost, but that growth has been negligible.

There has been a marked rise in cloud computing adoption among database teams, the survey found. Forty-one percent reported having cloud in production at scale or in limited use, up from 33% in the 2016 survey. Notably, 28% of respondents have cloud in production at scale, well over double that of 2 years ago (11%) (see Figure 1). One-third of respondents said their use of cloud is growing, with 20% reporting their cloud growth as "significant"—again, a rise over just 2 years ago (see Figure 2).

As public cloud adoption grows, much of this growth is driven by backup and recovery to support transaction environments. A total of 23% respondents delegate a significant

DATABASE EXPANSION

share (a quarter or more of their capabilities) of their backup and recovery processes for transactional environments to the public cloud. For additional processes affecting data environments, close to one in five rely on cloud for significant shares or their business continuity, monitoring, and provisioning processes. For analytical data environments, there is less commitment to public cloud at this time—at most, 17% of respondents are dedicating a notable portion of their backup and recovery process workloads to public cloud environments.

Cost reduction is the main benefit anticipated with cloud, but agility and capacity are more likely to be realized in existing deployments. Data managers and professionals seek the cost advantages of public cloud—which may form the basis of business cases, at least initially. Six in 10 foresee the cost reductions as the improvement sought with cloud computing.

As deployments mature, however, the additional agility and on-demand resources that the cloud brings to bear emerge as the leading benefits. In addition, the advantages of public cloud computing are far more apparent than 2 years ago. When asked about positives already seen, three in four said they have experienced greater agility. A majority cited the on-demand capacity clouds bring to the table. ■

Figure 1: Does your organization use any public cloud services for Oracle database and applications?

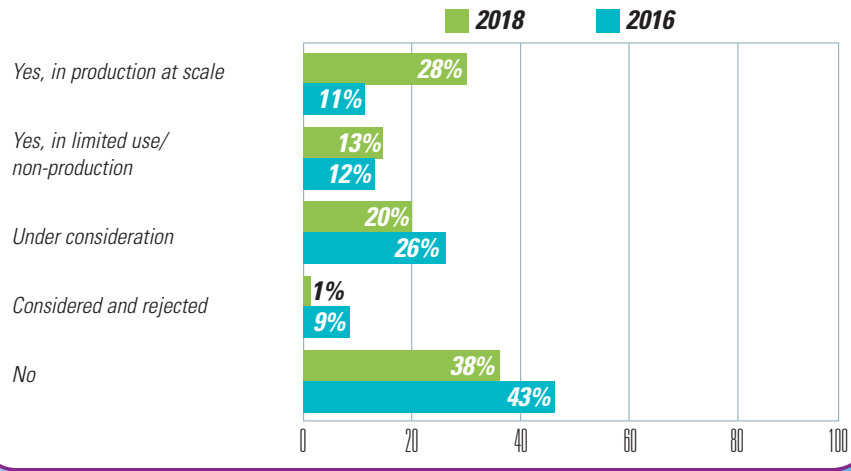

Figure 2: How has your use of public cloud services for Oracle database and applications changed over the past year?

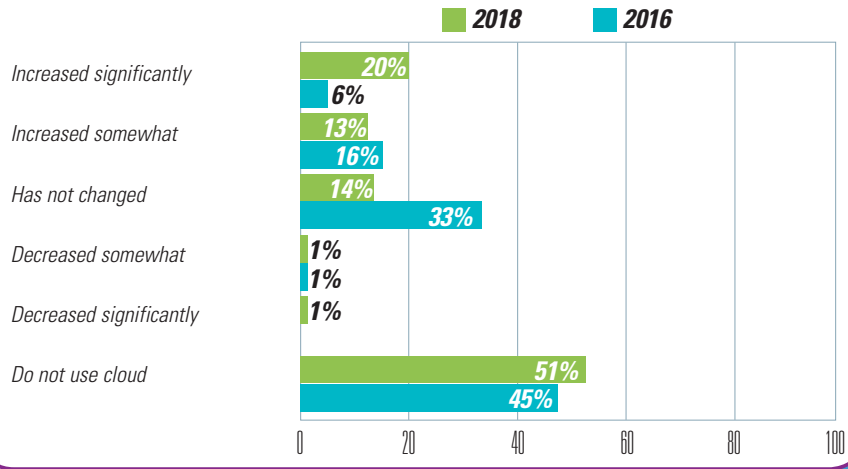


THE YEAR AHEAD
DATA
WILL DRIVE THE
ENTERPRISE
in
2019

BY JOE MCKENDRICK

The year just ending has been an interesting one for data managers. Artificial intelligence (AI) and machine learning took center stage, which also meant an increasingly glaring spotlight on data sourcing, management, and viability. The continued rise of the Internet of Things (IoT) also meant no letting up on demands for data environments to deliver requirements fast and furiously. The year ahead will bring more of the same—as well as a continuation of the transformation of information management.

Here are some of the changes and challenges on the horizon for 2019, as seen by leading industry participants and observers.

MASSIVELY DISTRIBUTED DATA GETS EVEN MORE DISTRIBUTED

Things keep moving away from the center. As Ken Tsai, global VP and head of cloud platform and data management for SAP, put it, customers' data now sits in an average of six to eight clouds, as well as their own data centers. "This is contributing to the rapid evolution of data processing technologies to become massively distributed," he said. He added that "data integration technologies are shifting from extract-transform-load to a process- and pipeline-driven approach, with data management and governance with the new sources. Applications and data platform technologies are, and will continuously get infused with, intelligent technologies so business scenarios can be automated by AI and heuristic usage data, creating a new class of solutions."

DATA GETS EVEN MORE STRATEGIC

In the year ahead, data managers will continue seeing their mandates extend well beyond their original course of action for managing and securing day-to-day transactions. It's now about "leveraging information to make

*In the year ahead, apps and data platforms will increasingly be infused with **intelligent technologies** so business scenarios can be automated by AI and **heuristic usage** data.*

capabilities to support both centralized and federated models."

At the same time, Tsai cautioned, this high level of distribution will "increase cost and complexity in managing mixed data landscapes and locations and it is difficult to know where the data resides, what data is available, as well as how to govern and trust the data source and accuracy, and monitor its usage and lineage across this distributed environment."

This will change the way information is managed, Tsai added. "Information won't be tethered to just one system and will flow freely and be connected—no matter how the business needs evolve or where the data consumer is located, or the device they want to use for access." Enterprises will approach information orchestration from raw feeds with intelligence and real-time analysis on vast quantities and varieties of data types, he said. This will "make it much easier to enrich the mission-critical information strategic, operational, and tactical decisions that result in increased revenues, improve operational efficiencies, and enhance customer experience," said Satya Sachdeva, VP of insights and data for Sogeti USA. Systems are evolving with this growing mandate. "Database technologies used for information management have been rapidly evolving from traditional relational database management systems and OLAP technologies to MPP-based appliances," he said. "Data lakes and Hadoop-based data environments for ingestion, wrangling, and analytics are gaining footholds across many organizations."

DATA GOVERNANCE GAINS

Data governance will continue to become even more critical in the year ahead. The EU's General Data Protection Regulation (GDPR) was a major driver for data governance programs in 2018, observed Emily Washington, senior vice president of product management at Infogix. "Suddenly, every organization conducting business in the European Union needed a well-functioning data governance program."

However, Washington continued, "this new focus on data governance didn't automatically translate to a better understanding of enterprise data and new analytical insights. Although many businesses are now collaborating across departments, there is still a disconnect between the analytics team and those that are focused on governance efforts, resulting in various metadata definitions across teams."

Introducing governance to analytical models will help businesses "aggregate the metadata around their models to ensure all teams have a complete understanding of their data and can leverage it in analytical models," she stated. "More businesses are now successfully cultivating open communication between various departments because of their data governance programs. This suggests that many companies are ready to expand their data governance program to a more strategic focus beyond just governing data."

MORE CLOUDS ON THE HORIZON

For some time, the idea of relying on public cloud services to handle critical or sensitive data assets was a non-starter. Attitudes have changed, and in the year ahead, public cloud may be the go-to solution at many enterprise data sites. "Public cloud solutions have made resource consumption cheaper, simpler, and dynamic," said Gaurav Yadav, founding engineer and product manager for Hedvig. There's another advantage beyond this, as public cloud services offer a degree of standardization that is needed across global enterprises or partner networks. "Businesses are struggling to keep up with this extreme pace of data generation and traditional data analytics tools are not capable of handling such globally distributed data."

While public cloud solutions have long been the preferred option for ▶▶

melissa

Multiplatform Data Quality Management—Unison is Speedy, Secure & Scalable

WHEN IT COMES to data quality, speed and security are paramount. We kept that in mind when building Unison, an ideal, multiplatform solution to establish and maintain contact data quality at higher speeds—processing 30+ million addresses per hour—while still meeting the most stringent security requirements.

Designed to be innovative, scalable and reliable, Unison is a holistic platform that sets the standard for data quality management. It brings data standardization, validation and enrichment together for speedy end-to-end data quality. With Unison, you can even cleanse your most sensitive customer information throughout the enterprise, as data never leaves your organization.

End-to-End Data Quality at Lightning Speed

Unison offers superior data quality management in one handy, on-premise platform. It was designed to save you time with simple setup, fast processing and automated updates.

Getting started requires no technical knowledge, coding or programming – and development time is completely eliminated. Simply install the web-based client-server application for full access to address, phone and email verification, plus geocoding capabilities.

The login portal is accessible through local intranets in any modern browser and is set up for effective project collaboration across multiple users. Unison's elegant interface is organized by the steps of the data management workflow and is fully customizable.

Once installed, Unison validates U.S. and Canadian addresses and performs aggressive address corrections at lightning speed. It leverages a Canada Post® SERP Certified™ and USPS® CASS Certified™ address engine to match and correct spellings and naming mistakes for cities and streets, and to add the correct street name suffix, prefix and ZIP+4 information.

Unison can also append latitude and longitude coordinates with Census data, validate and standardize email addresses and phone numbers, plus validate and parse full names in record time.

Data can be managed confidently on-site to meet compliance and security requirements due to Unison's ability to work completely offline. Even within the platform itself, Docker containerization security features leave assets isolated and self-contained, making Unison ideal for industries with sensitive data that must remain onsite.

Scalability & Flexibility

Unison supports horizontal scaling, so you can maximize existing hardware and any hardware you choose to add in the future. Scaling horizontally allows you to spread work across multiple servers and begets huge leaps in processing speed. Docker Swarm tears the roof off any performance ceiling by easily spreading Unison containers across as many processors as you choose to configure. The only limiting factor will be your network speed.

You also have the ability to connect to multiple RDBMS platforms and schedule jobs to process during off hours to maximize performance, then visualize analytics with reports generated automatically to better understand your data and collaborate on projects.

More than quick processing and setup, Unison saves you time by automatically distributing updates directly to your IT team. If your Unison installation has internet connectivity, updates will prepare themselves on your servers, and you'll get notified once it's ready.

Cleanse Sensitive Customer Info Securely & Safely

We've also added an additional layer of security so only authorized users can access the application. Administrators can choose to integrate Unison within the company's LDAP system for preexisting logins, or create account logins with role-based capabilities and configurable user rights through the Unison Authentication System.

Wide Compatibility for Future Enhancements

Unison currently supports Oracle, SQL Server, MySQL database, and a variety of delimited flat files to deliver address, email and phone verification, and geocoding for U.S. and Canadian records. In the future, Melissa will expand Unison's service offerings past contact data to include the full spectrum of data enrichment services for all data types. Streamline data prep workflows, reduce analytics busy work, gain more insights and increase efficiency with Unison.

Try a Unison product demo! Connect with members of Melissa's customer support team at www.melissa.com or call 1-800-MELISSA. ■

Melissa ■ 22382 Avenida Empresa ■ Rancho Santa Margarita, CA 92688-2112
Phone: 949-858-3000 ■ Fax: 949-589-5211 ■ Toll Free: 1-800-635-4772 ■ www.melissa.com

startups, existing organizations are also recognizing the advantages of such services. Most startups find it easy to ramp up by consuming public cloud solutions at first, and slowly migrating data back on-premise for cost optimization, Yadav said. However, more mature businesses always had data on-premise, and they migrated some of that data to the public cloud in order to reduce the number of resources needed to manage their infrastructure.

"Public cloud and 'as-a-service' solutions have transformed the expectations such that high-availability, simplicity, and security are expected to be built-in for consumer as well as enterprise products. It's not enough to just fulfill the customer requirements anymore, delighting the customer is the key," he noted

For the year ahead, expect to see "a huge push in data management vendors providing deeper public cloud integrations and removing the complexity out of hybrid cloud data management and analytics features," Yadav said.

Itamar Ankorion, CMO at Attunity, predicted greater use of major cloud infrastructure and platform providers—such as Amazon S3, Azure Data Lake Store (ADLS) and ADLS Gen2, and Google Cloud Storage—for cloud-based data storage, especially as organizations move to analytics-driven strategies. "These converged platforms host preferred analytics systems such as the Snowflake and Amazon Redshift data warehouses, and Amazon Athena," Ankorion pointed out. "These systems incorporate familiar SQL structures."

The traditional cloud benefits apply here, Ankorion continued. "Enterprises gain economic storage, processing and management, elastic resource consumption, and the shift of CAPEX to OPEX. They also can reuse storage for new analytics workloads. All this is changing how organizations manage data because they can more cost-effectively address more advanced analytics use cases and thereby realize a greater return on their data investments."

However, there will be a some challenges in achieving these cloud storage benefits, Ankorion cautioned. "Transferring high data volumes across wide area networks can result in high latency and consume costly bandwidth," he explained. "Enterprises also need to avoid lock-in and maintain full data mobility, moving workloads between clouds and across hybrid architectures based on lessons learned and changing requirements."

*Data governance will continue to become even **more critical** in the year ahead.*

IT'S HYBRID TIME

Many enterprises have on-premise data assets they want to keep within their four walls, at least for now. With the rising popularity of public cloud, "2019 will be the year of multi-cloud and hybrid cloud," said Arun Murthy, CPO and co-founder of Hortonworks, which recently announced a merger with Cloudera. "Cloud providers will, more aggressively, differentiate among each other in specific areas—such as operational readiness at Amazon, enterprise integrations for Microsoft, and AI and machine learning at Google," he predicted. This array of capabilities from across vendors will give rise to multi-cloud strategies being put into motion, bound by "common security, governance, and data or workload management strategies. We will also start to see enterprises move some always-on workloads from the cloud back on-prem for a hybrid model to optimize economics."

STRENGTHENING THE DATA FOUNDATION

As data becomes the lifeblood of most business models and strategies, expect to see a boom in the use of algorithms throughout the enterprise for a variety of purposes. "This has the highest opportunity for business payoff, especially for augmenting human intelligence," said Suketu Gandhi, partner in the digital transformation practice at A.T. Kearney. He predicted that as a result, "data and organizational intelligence will be discussed at the board level, and CEOs will be talking about it regularly as well." Gandhi cautioned, however, that data managers need to get ready for the push to an algorithm-driven enterprise. "It shines a light on the state of data—both internal and external—and the ability of an enterprise to use it. There is a lot of work to get ready for the fundamental reality of structured versus unstructured, internal versus external, real time versus batch," he added.

Business and data executives are conducting reality checks on the best approaches for the employment of algorithms, autonomous operations, and AI. "Many businesses are being smart when it comes to adopting new technologies," said Lyndsay Wise, director of market intelligence at Information Builders. "Instead of buying into the hype, they are asking critical questions for garnering the strongest ROI, resulting in a delay in broad adoption. For instance, organizations are realizing that strong data management is a core foundation for predictive and AI technology and are first focusing efforts on getting their data house in order. Others have realized that they don't have the pool of data necessary to make the most of predictive technologies and are investing in building the right data streams."

Most organizations by now "have worked with data long enough to know when they're ready for a new trend or significant investment," Wise added. "Organizations that are moving forward with predictive analytics, machine learning, and AI are doing so because they've dedicated enough time and resources to their data management and are confident they have the right amount of data needed. Data comfort is really critical here. You can't implement data-fueled technologies if, as an organization, you're not comfortable with data internally."

database TRENDS AND APPLICATIONS
The Journal of Information Integration and Management

Technical Community College Achieves Smarter Data Integration with Kore Technologies

Kourier Integrator bridges the gap between the Colleague ERP system and custom applications.

Serving the citizens of Forsyth and Stokes counties in North Carolina since 1960, Forsyth Technical Community College offers vocational instruction and training in skilled trades as well as college transfer and two-year degree programs, corporate training, continuing education, and personal enrichment classes.

THE CHALLENGE

Forsyth Tech has relied on Colleague by Ellucian on Rocket Software's UniData (MultiValue) platform as its central ERP system for years. The Colleague system is the backbone of the College, handling everything from an administrative perspective, including student information, registration, finance, and human resources (HR).

Over time, however, Forsyth Tech found the need for new capabilities, and purchased additional applications. The College also developed its own custom applications, which is less expensive and also means applications meet its own business rules. Integrating data within the core ERP system with those applications, however, proved to be challenging.

While there are plenty of solutions available for relational systems, tools for reporting and integration for MultiValue systems are not as plentiful and, in many cases, cost-prohibitive. Moreover, as a state-funded institution, Forsyth Tech puts a high priority on cost-effectiveness in everything it does in order to concentrate its resources on its primary goal of educating, engaging, and retaining students.

THE SOLUTION

Forsyth Tech initially became familiar with Kore Technologies through a data warehousing and reporting project it embarked on in summer 2017 using Kore's Kourier Integrator, SQL Accelerator and Operational Data Store for Colleague. However, as the Forsyth Tech team members evaluated more of Kore's product suite, including Kore's RESTful (Representational State Transfer) Web Services and REST Gateway, they realized that this would be ideal for accessing data in real time from the Colleague ERP system to their applications. Integration quickly became a key focus.

"All the tools were there from a data warehousing perspective," said Christopher Pearce, associate vice president and CIO of Forsyth Tech. "When we looked at the cost and ROI, plus the add-ons the College was considering, we were impressed. It was everything that the College was looking for at a price point that the College could commit to easily. Once the College found out about the REST tools, we knew that we had to start working on the relationship with Kore."

> **When we looked at the cost and ROI, we were impressed.**

THE BENEFIT

Forsyth Tech's custom apps handle course substitution, employee evaluation, and tutor requests, and now, because of Kore, they can all be updated in real time. "Up until this point," said Pearce, "data transfer from Colleague to its custom applications was accomplished using rudimentary tools to get data into a flat-file format, and then the data was passed over to the application, which had to import the file before the community could start using the data. The process was a 'management nightmare' that required writing scripts and scheduling data transfers, and had a level of latency that is unacceptable by today's standards."

Kore's Kourier REST solution accelerated the process of building secure, real-time REST APIs for the integration. Beyond real-time integration with its current apps, Forsyth Tech has also been able to better serve the needs of additional college departments and expand on its initial set of in-house applications. "For example," said Pearce, "each year, the HR department updates the organizational chart to identify who works where and to whom they report. Building on the data integration with the employee evaluation application, Forsyth Tech was able to create a new interactive organizational chart in just a few days, as opposed to months. In addition, there are now plans for new projects, including bi-directional data integration for its custom applications as well as integration between its Colleague ERP system and third-party applications."

NEXT STEPS

Next on the roadmap is a project to tackle better understanding of and reporting on student attendance. A new attendance application will allow the College to provide improved attendance information to the state, which is tied to its funding. In addition, the application will enable easy access to data for use with predictive analytics tools to quickly identify students who may be struggling academically so they can be supported with tutors or counseling.

Reflecting on its use of Kore's solutions, Pearce said, "the advantages go beyond data integration. With the use of new approaches like REST, Forsyth Tech is achieving its goals of persistence, engagement, and retention by enabling greater access to information for both staff and students to support academic success. It is all about giving students the information they need when they need it and removing barriers to furthering their education." ■

CONTACT INFORMATION

kore Technologies

5186 Carroll Canyon Road, Suite B
San Diego, CA 92121
866-763-5673
www.koretech.com
info@koretech.com

*A **stronger data foundation** will help deliver the value expected from data-intensive **methodologies such as** predictive analytics, machine learning, and AI.*

A stronger data foundation will help deliver the value expected from data-intensive methodologies such as predictive analytics, machine learning, and AI, said Wise. "When implemented correctly, the opportunities are limitless. We can't really predict the full scope of how AI and predictive technologies will improve organizations, but we know there will be significant improvements in efficiencies, cost savings, customer service and experience, bias improvement, and employee growth—to name a few."

ECOSYSTEMS EVOLVE

Success in the data-driven organization means more than relying on corporate databases—innovation and growth now are tied to the development of connected ecosystems of partners, customers, and other constituencies. Expect to see a greater focus on these networks in the coming year. "Organizations already have started to pull the camera back and take a broader look at how their ecosystems drive their businesses, but it's about something more than connecting or integrating," said Frank Kenney, director of sales enablement at Cleo and former Gartner analyst. "It's about driving value. In the next year and beyond, savvy organizations will deliberately focus on enabling their ecosystems to obtain different types of value out of relationships with customers, partners, partners' customers, and so on. That means they are integrating not only their dynamic networks of people, partners, customers, systems, applications, and things but also the processes and interactions that drive those relationships."

Managing and capitalizing on ecosystems "requires a modern approach to integration, one that can enable your business to take advantage—through functionality, through interoperability, through visibility—of all the ways value can be created in today's era of digital transformation," Kenney said.

Organizations aren't quite ready to fully embrace an ecosystem-driven approach, however. "More work needs to be done in how we view and manage all the data interactions happening in our ecosystem, but it's achievable," Kenney said. "It's important to set realistic goals and expectations, and simply mapping out how your intelligence network has expanded is a great first step to learning to use the data."

MORE TARGETED SEARCH

Enterprise search is an area that will increasingly become part of organizational information management strategies in the year ahead, especially as both data and content are embraced as strategic assets enabling business growth. Jill Shuman, director of project engagement of the Copyright Clearance Center, foresees the rise of "a search solution that gathers information from multiple, disparate content sources all at once and presents results on a single page."

Ideally, organizations "should consider creating a consolidated place to store both internal and external content, coupled with a single enterprise-wide search function," said Shuman. "This allows employees access to everything they need in one basic search effort instead of having to check each disparate source individually. This approach saves time, money, and most importantly, protects organizations from loss of institutional knowledge, which could cost millions of dollars."

As consolidated enterprise search emerges, companies will see "more productive staff, because they won't have to duplicate the efforts of past employees—the information would be captured before the previous employee left and be easily available," said Shuman. "An employee who is fully engaged with a company and its information feels more a part of the team."

AI technologies are also playing a key role in enhancing both data and content search. "Coupling technologies such as machine learning and natural language processing with search is helping companies to glean greater value and insight from all their data, both structured and unstructured," said Kamran Khan, a managing director with Accenture Applied Intelligence. "It increases the ability to find, analyze, understand, and present data more accurately and efficiently and provides a new turbo-charged way of searching that helps companies not only gain more insight from their data but it also improves the user or customer experiences that directly support many business goals and objectives. Examples of these might be an intranet or KM application, ecommerce site, or customer service portal."

However, AI-driven information technologies are still immature, Khan cautioned. "While technical leadership might understand machine learning and NLP, there is still a knowledge gap with regard to fully realizing how these technologies work and how they can play a role in improving search and content analytics. These technologies are also not easy to implement. It takes time, money, and a deep understanding of content analytics to bring true cognitive search into enterprises."

AI technologies are changing the way information is managed, said Khan. "Applying NLP and machine learning to search applications, for a smarter way of searching and understanding data will absolutely change the way information is managed. Right now, nearly 80% of all enterprise data is unstructured content which is rarely used in analytics. This number will only continue to grow as each enterprise's data continues to increase." ■

Moving to a MODERN DATA ARCHITECTURE

database
TRENDS AND APPLICATIONS

Best Practices Series

9 STEPS FOR MOVING TO A MODERN DATA ARCHITECTURE

Best Practices Series

GETTING TO A MODERN data architecture is a long-term journey that involves many moving parts. Many organizations have vintage relational database management systems that perform as required, with regular tweaking and upgrades. However, to meet the needs of a fast-changing business environment, data executives, DBAs, and analysts need to either build upon that, or re-evaluate whether their data architecture is structured to support and grow with their executive leaderships' ambitions for the digital economy.

In terms of technology, a combination of cloud resources and automation is helping data executives keep up with emerging requirements—artificial intelligence (AI), the Internet of Things, and real-time customer service. A recent survey of 260 data managers, fielded by Unisphere Research, a division of Information Today, Inc., in partnership with VMware and the Independent Oracle Users Group, found a noticeable rise in cloud computing adoption among database teams. Forty-one percent had cloud in production at scale or in limited use, up from 33% in a similar survey conducted in 2016. In addition, one-third said their use of cloud is growing, with 20% describing their cloud growth as "significant"—again, a rise over just 2 years ago. Many of these deployments are in pursuit of greater automation to deliver backup and recovery requirements, a vital piece of all data architectures.

In addition, there has been notable growth in adoption of new systems and frameworks that are helping to shape the foundation of modern data architectures, according to a Unisphere Research/Oracle survey of more than 200 companies. Adoption of data lakes—a place to store diverse datasets without having to build a model first—continues to rise as data managers seek to develop ways to rapidly capture and store data from a multitude of sources in various formats. Overall, 38% of organizations are employing data lakes as part of their data architecture, up from 20% in the 2016 survey. In addition, Apache Spark, an open source analytics engine targeted at large-scale data processing at real-time speeds, is being used at one in four organizations as part of their data architectures, with another 15% considering adoption. One

in four executives are also employing machine learning, a part of AI.

There are many promising technologies converging to support the development of a modern data architecture, but organizations need to be prepared and ready to capitalize on such new capabilities. Here are ways to prepare enterprises for this journey:

Define what a modern data architecture means to the business.

A "modern data architecture" can mean many things, and implementations will vary from organization to organization. For some, building and supporting big data with Apache Hadoop clusters is an important element of a modern data architecture. Others may have even moved beyond Hadoop and are creating systems that respond to real-time events. To many enterprises these days, a modern data architecture means shifting data management functions to the cloud.

Make data as widely available as possible to those who need it.

Analysts, pundits, and vendors have been talking about "data democracy" within organizations for close to 2 decades; the technology and know-how are finally available to make this possible. A data analytics-driven enterprise needs to support informed decisioning all the way down the corporate ranks—for example, call center representatives need relevant insights about the customers with whom they are interacting, and production workers need access to data on machine performance.

Build a self-service culture.

Self-service capabilities are the key to data democracy. By putting analyses and reporting in the hands of end users who need it when they need it, data-driven thinking will become a standard part of day-to-day business, not the domain of quants and upper-echelon executives. Key to a self-service culture is effective data governance to not only provide the "guardrails" of data integration and consumption, but to also set the course for the business.

Get real about real time.

Real-time responsiveness to customer demands, as well as end-user requirements for data, is what now separates the disrupters from the disrupted in many markets. Real-time capabilities are reaching commodity status with built-in in-memory processing in major vendor products as well as open source frameworks such as Apache Spark. A modern data architecture needs to be able to support real-time interactions and calls for data where and when it is needed.

Make data analytics part of your digital transformation.

Successful digital transformation—now being sought by organizations across the board—depends on the ability to capture and analyze data streaming to and from customers, partners, and employees. Whether it is business intelligence or AI, any digital transformation strategy requires a modern data architecture to succeed.

Bring data closer to the customer.

Designers of a modern data architecture need to ask one simple question: What does the customer need? Examining what customers—as well as internal end users—require in terms of data and insights helps pave the way to a responsive model and approach to data delivery. Exploring customer and user requirements isn't limited to the first rollout of the architecture—data architecture proponents need to keep asking if their environments are continuing to meet customer and end-user needs, and what needs to be changed.

Automate as much as possible.

A modern data architecture has to be highly scalable, capable of not only streaming and ingesting massive amounts of data but also tagging the data for analytics algorithms and AI datasets.

Encourage a collaborative approach.

Until recent times, data environments were the domain of data executives and DBAs. With data now an essential component of business strategy, managers across the enterprise need to have input on how data is discovered, streamed, and analyzed. While IT and data managers will still oversee the backend functions of systems—such as security and availability—business end users need to be able to provide their input into how data is made available.

Take a data platform approach.

Platform thinking, in which participation from across the ecosystem is encouraged and digital resources are made available, can open up data and insights to an ecosystem of partners, thereby building on a network effect. This maps to the rise of digital enterprises, in which interactions and transactions take place across a platform, versus traditional sales channels. This so-called plug-and-play business model enables data producers and data consumers to develop and engage with services connected with data analytics.

A modern data architecture is essential to move forward in today's digital economy. While the technology is available and affordable, organizations need to be ready to take full advantage of a more responsive, real-time, analytical environment. ■

—Joe McKendrick

DELPHIX

Why Every Enterprise Organization Needs DataOps

Data is the new strategic asset in today's economy. It is a source of unique insights and the engine of modern applications and machine learning. Winning in this new software-driven landscape requires aligning people, process, and technology around the secure flow of data to accelerate innovation and drive differentiation. The alignment hinges on the easy access to any data in any location, without compromising data privacy or compliance.

As enterprises become more data-driven, the explosion of data demand can sometimes create intense friction with the ever-increasing cost, complexity, and risk of managing, distributing, and securing that data. The accelerated adoption of different types of databases, like MySQL, MongoDB, PostgreSQL, Amazon RDS, and AWS Aurora, in addition to traditional databases, such as Oracle, and Microsoft SQL, increases this friction manifold.

Data friction can be solved with DataOps. Like DevOps, DataOps is a people problem, with technology as a key enabler of data flow and reliever of this friction. DataOps, as seen with a few early adopters, plays an important role within the enterprise to have an immediate impact on the following:

- Software Development—Automated, repeatable, and faster provisioning of higher quality data for AppDev, QA, and test environments that results in higher velocity for new releases and feature updates.
- Data Privacy and Governance—Identifying and mitigating risk to meet compliance and privacy constraints associated with data flow between data managers, like DBAs, and data consumers, like developers.

The key to DataOps success lies in a technology platform that provides the following three capabilities:

- Enterprise-Readiness—A platform that seamlessly integrates and scales with the heterogeneous enterprise IT landscape and works with all relevant data wherever it exists. No point solutions to stitch together to implement a data strategy.
- Frictionless Data Delivery—Data as a self-service, when and where it is needed. An end-to-end platform captures the data from where it exists, moves it to where it needs to be, prepares it for relevant use cases, such as data pods for SDLC, and lastly, presents masked data to data consumers on a day-to-day basis with minimal overhead or delay.
- Integrated Risk Management—The platform proactively identifies risk with de-identification, obfuscation, and other techniques to mitigate risk as data flows across the enterprise.

DevOps success has come from breaking down silos between teams, bringing speed and agility through close collaboration and automation via technology platforms. Delphix addresses similar technology aspects for DataOps. Data remains secure, moves fast, and provisions instantly, which helps organizations to do development and shift-left testing faster. This results in higher quality software and the ability to release more features. With data pods, developers, testers, and CI/CD pipelines can easily provision and refresh realistic masked production data in a matter of minutes. Robust branching and bookmarking capabilities can be used to create a readily available test data catalog of production, synthetic, and curated datasets.

Let's take Channel 4 Television as an example to illustrate how it leverages a technology platform called Delphix to adopt a DataOps approach and drive business success.

Channel 4 Television, a UK-based, multibillion-dollar media company, delivers a high volume of programming and commercial content to millions of viewers every day. It relies heavily on a suite of data warehouses to provide accurate and up-to-date insights into business performance and future opportunities. Their ability to rapidly refresh this information and manage volume sizes was key to their ongoing success and helped target their spend on the screen instead of on technology. They already used Agile and DevOps models to accelerate delivery of business change, but a lack of automation tools to optimize data management caused delays to environment deployment into several weeks and drove significant costs in storage and tooling. These delays caused the team to miss capitalizing on several market trends and burdened them with additional capital expenditure. The protracted time taken to refresh the circa 18TB of data in test environments inhibited system and business agility, ultimately preventing teams from realizing their goal of regular change and deployment.

Channel 4 now uses Delphix as part of a DataOps approach to optimize and automate their application development process. Changes and modifications to key systems are more efficiently managed and the availability of reporting data is improved. The Delphix platform has also enabled the acceleration of data delivery for development, testing, and reporting to drive faster innovation and improve operational efficiency. Finally, data pods have improved the connection between DBAs to developers to ensure fresh, secure delivery of data. What used to take days now takes just minutes. Today, Channel 4 is able to release applications quickly and with more confidence using the Delphix Dynamic Data Platform. ■

DELPHIX
www.delphix.com

AEROSPIKE

Maximize the Value of Your Operational Data

Modernize and Transform Your Enterprise Via Real-time Transaction and Analysis Processing

OVERVIEW

It might sound too good to be true: a database system that processes large volumes of operational data in real time while delivering exceptional runtime performance, high availability, and cost efficiency while still keeping your data safe. What if early adopters in banking, telecommunications, and other industries are already harnessing such a database for achieving results that are transforming their businesses in myriad ways? What if published benchmarks demonstrate sub-millisecond response times for high throughput read/write workloads over high data volumes with substantial cost savings compared with traditional alternatives?

This paper introduces key technologies that Aerospike clients are using to modernize their data management infrastructures and realize such impressive (and seemingly impossible) results as:

- Rapid read/write speeds without extensive tuning or a separate data cache
- Substantially smaller footprints than popular alternatives, often leading to 3-year total cost of ownership (TCO) savings of $3-5 million per application
- 24x7 availability, including cross-datacenter replication
- Operational ease during scale-out and maintenance
- Interoperation with popular software offerings, including Apache Hadoop, Spark, and Kafka

Sounds unbelievable, right?

THE TECHNOLOGY IN BRIEF

Aerospike provides a distributed, highly scalable database management system for demanding read/write workloads involving operational data. It was designed to deliver extremely fast—and predictable—response times for accessing data sets that span billions of records in databases of 10s – 100s TB. Other design features address fault tolerance and near 100% uptime even during upgrades and maintenance.

How? By capitalizing on proven architectural approaches—such as distributed computing and parallelism—and developing new technologies to meet business demands that hadn't even surfaced when older systems were originally built. Indeed, Aerospike's patented Hybrid Memory Architecture™ (HMA) drastically reduces traditional I/O and network communication compared with other approaches; it also uses CPU resources considerably more efficiently. The cumulative impact of these features (and others) enables Aerospike to deliver remarkable speed at scale.

APPLICATIONS AND USE CASES

Applications that benefit from Aerospike typically share some or all of these characteristics:

- Service-level agreements (SLAs) that require sub-millisecond database response times
- High throughput for mixed workloads (e.g., 3–5 million operations per second)
- Support for managing billions of business records in databases of 10s–100s TB
- High availability and fault tolerance for mission-critical applications
- High scalability for handling unpredictable increases in data volumes and transactions
- Adaptable infrastructure for managing varying types of data with minimal effort
- Low total cost of ownership (TCO)

KEY FEATURES AND TECHNOLOGIES

Aerospike is a shared-nothing database system that operates on a cluster of commodity server nodes:

- It's a schema-free, key-value data store.
- Aerospike exploits volatile and non-volatile memory in a distinctive way, providing rapid access to index and user data.
- An intelligent client layer minimizes costly network "hops" needed to access data.
- Immediate record-level consistency and high availability are guiding principles.
- Access management controls and transport encryption protect sensitive data.
- Asynchronous replication across data centers provides disaster recovery.
- Ready-made connectors, a publish/subscribe messaging system, and partner offerings help firms integrate Aerospike into their existing IT infrastructures. ■

FULL REPORT

To get a copy of the full report, please go to: www.aerospike.com/maximize-operational-data.

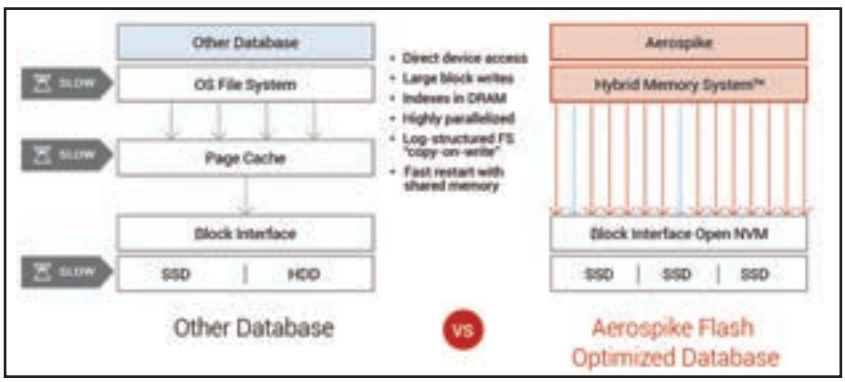


TREND-SETTING
PRODUCTS for
2019

You can call it the new oil, or even the new electricity, but however it is described, it's clear that data is now recognized as an essential fuel flowing through organizations and enabling never-before-seen opportunities. However, data cannot simply be collected; it must be handled with care in order to fulfill the promise of faster, smarter decision making.

More than ever, it is critical to have the right tools for the job. Leading IT vendors are coming forward to help customers address the data-driven possibilities by improving self-service access, real-time insights, governance and security, collaboration, high availability, and more.

To help showcase these innovative products and services each year, *Database Trends and Applications* magazine looks for offerings that promise to help organizations derive greater benefit from their data, make decisions faster, and work smarter and more securely.

This year, our list includes newer approaches leveraging artificial intelligence (AI), machine learning, and automation as well as products in more established categories such as relational and NoSQL database management, MultiValue, performance management, analytics, and data governance.

On the following pages, we present *DBTA*'s list of Trend-Setting Products for 2019. We encourage you to continue your exploration by visiting the companies' websites for additional information. In addition, in this issue, we include Product Spotlight articles penned by company executives with explanations of what makes these products unique.

And, to stay on top of the latest news, IT trends, and research, go to www.dbta.com, and review *DBTA*'s extensive collection of white papers and research reports at www.dbta.com/DBTA-Downloads/White Papers. ■

ACTIAN

www.actian.com

Actian Zen Core Database for Android—a new member of the Zen Enterprise Edge Database family that focuses squarely on the needs of edge and IoT application developers, providing persistent local and distributed data across intelligent applications embedded in smart devices

AEROSPIKE

www.aerospike.com

Aerospike—an enterprise-class, NoSQL database solution for real-time operational applications, Aerospike is offered in two versions: the open source Aerospike Community Edition and the Aerospike Enterprise Edition, which includes all the features of the Community Edition, plus additional premium features

ALTR

www.altr.com

ALTR Data Security Platform—having emerged after nearly 4 years of stealth, the ALTR platform is designed to eliminate threats to data through core approaches in private blockchain, in-line data techniques, real-time alerting, and reporting for business enablement, and offers support for all major database platforms

AMAZON WEB SERVICES

https://aws.amazon.com

Amazon Aurora Serverless—bringing the power of the MySQL-compatible database built for the cloud to applications with intermittent or cyclical usage patterns, Aurora Serverless is a recently added deployment option that automatically starts, scales, and shuts down database capacity with per-second billing for applications

ARCADIA DATA

www.arcadiadata.com

Arcadia Enterprise—with a data-native and AI-driven architecture designed for big data, Arcadia Enterprise improves access to analytics and BI for business analysts, enabling them to explore data with recommendations that visualize and identify insights, and to derive value with minimal IT overhead

ATTIVIO

www.attivio.com

Attivio Cognitive Search and Insight Platform—understanding that the key to transforming the productivity of information workers is improving the relevancy of answers to search queries, Attivio combines natural language processing, machine learning, and knowledge graphing for a contextualized search and discovery experience with a secure enterprise platform

ATTUNITY

www.attunity.com

Attunity Replicate—helps to accelerate data integration for analytics with an intuitive GUI that eliminates the need for manual coding, and empowers organizations to improve data replication, ingest, and streaming across a range of heterogeneous databases, data warehouses, and big data platforms

BACKOFFICE ASSOCIATES

www.boaweb.com

Information Governance Cloud—a cloud-based SaaS solution that defines operational policies and business rules for data, aligns them with elements of the business strategy, and orchestrates their enforcement through any system, process, and stewardship platform—including the Data Stewardship Platform by BackOffice Associates

BLUEDATA

www.bluedata.com

BlueData EPIC (Elastic Private Instant Clusters) software platform—leverages the power of containers to help make it easier, faster, and more cost-effective to deploy big data analytics and machine learning—whether on premise, in the cloud, or in a hybrid architecture

BLUE MEDORA

https://bluemedora.com

BindPlane—offering IT monitoring integration as a service and a standardized approach to enterprise monitoring integration, BindPlane delivers Dimensional Data, which enhances health and performance metrics with critical metadata about the inter- and intrarelationships between the metrics in the IT stack

BRADMARK TECHNOLOGIES

www.bradmark.com

Bradmark's Surveillance DB—gives IT professionals the necessary visibility of overall system health, and offers cross-functional capabilities, including real-time monitoring, proactive event management, flashback/time-slicing, alarming and alerting, and SQL performance analysis with support across SAP Sybase, MS SQL Server, Oracle, Informix, and IBM DB2 databases

CAMBRIDGE SEMANTICS

www.cambridgesemantics.com

AnzoGraph—a native, massively parallel processing, distributed graph database that provides advanced analytics at big data scale and is available on a standalone basis for use behind the firewall as well as in the cloud

CAZENA

www.cazena.com

Cazena's Data Lake as a Service—with the goal of accelerating enterprises' ability to leverage analytics, machine learning, and artificial intelligence without hiring or re-training their workforces, Cazena's Data Lake as a Service, its flagship solution, offers a SaaS-like experience for analytics and business teams, with enterprise features for IT

CISCO SYSTEMS

www.cisco.com

Cisco Unified Computing System (UCS)—an integrated, scalable, multi-chassis platform for data center environments that combines compute, network, storage access, and virtualization into a system designed to reduce total cost of ownership and increase business agility and IT staff productivity

CLEARDB

https://w2.cleardb.net

ClearDB's MySQL Database-as-a-Service (DBaaS) solution—delivers a geographically redundant database infrastructure to help ensure that continuous access to a customer's MySQL database assets in the cloud is not only resilient within local clusters but across geographies

CLOUDERA

www.cloudera.com

Cloudera Data Science Workbench—with Python, R, and Scala directly in the web browser, the platform for collaborative data science delivers a self-service experience with connectivity not only to CDH but also to the systems data science teams rely on for analysis

COGNITIVESCALE

www.cognitivescale.com

Cortex5 software—simplifies the design, development, delivery, and management of enterprise-grade artificial intelligence systems that weave knowledge and learning across the enterprise, and helps businesses apply artificial intelligence and blockchain technology to solve complex business problems in the financial services, healthcare, and digital commerce markets

COLLIBRA

www.collibra.com

Collibra—designed to address the gamut of data stewardship, governance, and management needs of complex, data-intensive industries, the Collibra configurable cloud-based or on-premise solution puts people and processes first—automating data governance and management to quickly and securely deliver trusted data

COUCHBASE

www.couchbase.com

Couchbase Server—built on NoSQL technology, the database software platform delivers performance at scale, in any cloud, and features such as a memory-first architecture, geo-distributed deployments, and workload isolation, plus a SQL-compatible query language (N1QL) to make migration from RDBMS to Couchbase Server easier

DATA DYNAMICS

www.datadynamicsinc.com

StorageX Dynamic File Management—helping to free data from technology lock-in, complexity, and risk, the platform provides enterprises with an intelligent, policy-based, cloud storage management platform to enable data portability, usability, and insight for business agility and operational efficiency

DATABRICKS
https://databricks.com

Databricks' Unified Analytics Platform—a cloud-based platform powered by Apache Spark that enables data science teams to collaborate with data engineering and lines of business to build data products; recently added new capabilities to lower the barrier for enterprises to innovate with AI

DATASTAX
www.datastax.com

DataStax Enterprise—powering the "Right-Now Enterprise" with the always-on, distributed cloud database built on Apache Cassandra and designed for hybrid cloud, DataStax Enterprise helps companies to exceed expectations through consumer and enterprise applications that provide responsive and meaningful engagement

DATAVAIL
www.datavail.com

Datavail Managed Services—helping clients to collect, manage, and derive value from organizational data, as well as streamline processes through software integration and custom development, Datavail Managed Services are provided by a team of more than 1,000 data professionals including DBAs, developers, project managers, consultants, and business experts

DATAWATCH
www.datawatch.com

Datawatch Angoss—a suite of big data analytics software solutions and consulting services aimed at providing customers with a competitive advantage by helping them to turn information into actionable business decisions across logistics, risk, marketing, and sales

DBI SOFTWARE
www.dbisoftware.com

pureFeat V7 Performance Management Suite for IBM Db2 LUW—solutions for database performance analytics, tuning, and trending; response time analytics and SLA attainment, alerting and real time monitoring, along with new features to help DBAs and management teams to cope with the volatility of agile environments

DELL EMC
www.dellemc.com

Dell EMC PowerEdge MX— high-performance, modular infrastructure that is designed for the software-defined data center and able to support a variety of traditional and emerging data center workloads, including dense virtualization, software-defined storage, software-defined networking, artificial intelligence, and big data projects

DELPHIX
www.delphix.com

Delphix Dyamic Data Platform—helps organizations provision lightweight, compressed copies of production data in minutes, while keeping them in sync; secure sensitive data in adherence with security policies; and move and manage data from any environment—on premise, cloud, or hybrid

DENODO
www.denodo.com

Denodo Platform—offers the benefits of data virtualization, such as the ability to provide real-time access to integrated data across an organization's diverse data sources without replicating any data, and offers broad access to structured and unstructured data residing in enterprise, big data, and cloud sources

DIAMANTI
https://diamanti.com

Diamanti's D10 bare-metal container platform—integrates open source Docker and Kubernetes with purpose-built hardware and container-granular control to provide a full container stack to give infrastructure architects, IT operations, and application owners the speed, simplicity, efficiency, and control they need to run stateful containerized applications at scale

DREMIO
www.dremio.com

Dremio Data-as-a-Service Platform—connects any analytical process to any data source and scales from one to 1,000 plus servers, running in the cloud, in Kubernetes, or in a Hadoop cluster to help analysts and data scientists work together to discover, curate, and collaborate for diverse analytical use cases

EMPOLIS INFORMATION MANAGEMENT

www.empolis.com

Empolis Service Express—a central knowledge platform offered as SaaS that speaks the language of its users and can be used by customers in self-service, first-level service staff, and experts, giving them direct access to all relevant information by entering a search query

ERWIN

https://erwin.com

erwin EDGE platform—with metadata connectors and automated code generation, sensitive data discovery, data mapping, and cataloging tools, erwin EDGE creates an "enterprise data governance experience" that facilitates collaboration between both IT and the business to unlock the value of data both at rest and in motion

FRANZ

https://allegrograph.com

AllegroGraph—employing a combination of semantic and graph technologies that process data with contextual and conceptual intelligence, the database technology enables businesses to extract sophisticated decision insights and predictive analytics from their highly complex, distributed data that can't be answered with conventional databases

GOOGLE

https://cloud.google.com

Cloud SQL—a fully managed database service hosted on the Google Cloud Platform that offers high performance, scalability, and convenience, helping to make it easy for users to set up, maintain, manage, and administer relational PostgreSQL and MySQL databases in the cloud

GRIDDABLE

https://griddable.io

Griddable.io—provides a SaaS platform for hybrid cloud data integration that synchronizes data across any topology or database platform on any cloud, and includes a policy engine that controls a flexible definition of what data is filtered, masked, or transformed in transit

GRIDGAIN SYSTEMS

www.gridgain.com

GridGain In-Memory Computing Platform—sits between the application and data layers to provide in-memory speed and massive scalability to applications built on disk-based databases, and works seamlessly with existing application and data layers including all popular RDBMS, NoSQL, and Hadoop databases

HEWLETT PACKARD ENTEPRISE (HPE)

www.hpe.com

HPE Digital Prescriptive Maintenance Services—a series of AI-enabled industry offerings from HPE Pointnext, that enable problem prevention and increase productivity of industrial equipment—predicting, suggesting, and automating the right action to fix a problem before it causes harm

HVR

www.hvr-software.com

HVR Platform—a scalable enterprise data integration and validation platform to simplify high-volume real-time data movement that includes built-in efficiency features such as log-based change data capture, data compression, as well as security features such as encryption and KMS support for AWS

IBM

www.ibm.com

IBM Cloud Private—a key element of IBM's hybrid cloud strategy that—using container technology—can be installed on a wide range of enterprise systems to create a private cloud with architecture and capabilities consistent with the public IBM Cloud

IDERA

www.idera.com

SQL Diagnostic Manager—provides monitoring and diagnostics that allow DBAs to detect database performance problems and their root causes quickly, then tune database health and performance—and recently introduced a SQL Query Tuner for optimizing SQL queries and indexes for SQL Server, and new MySQL support

IMANISDATA
www.imanisdata.com

Imanis Data Management Platform—big data enterprise data management software for backup, recovery, disaster recovery, test/dev, and archiving—that is architected for petabyte scale, is data-aware, and powered by machine learning to reduce cost, minimize risk, and enable insight to all of data

INFLUXDATA
www.influxdata.com

InfluxData Platform—a complete time series platform built specifically for metrics, events, and other time-based data—whether the data comes from humans, sensors, or machines—that helps developers build next-generation monitoring, analytics, and IoT applications faster and more easily

INFORMATICA
www.informatica.com

Informatica Intelligent Data Lake Management—products certified to run on popular, on-premise Hadoop distributions or cloud deployment enable data scientists and business analysts to quickly find the data they're looking for on a self-service basis using semantic and faceted search, while automatically understanding data lineage and data relationships

INFORMATION BUILDERS
www.informationbuilders.com

WebFOCUS—combines traditional governed BI with business-led agile analytics, enabling enterprises to deliver a broad range of reports, dashboards, documents, and applications while empowering business users to create their own data visualizations, charts, graphs, and infographics

INNOVATIVE ROUTINES INTERNATIONAL (IRI)
www.iri.com

IRI DarkShield—consolidates and multi-threads the search, remediation, extraction, and reporting of sensitive data in unstructured files sitting in multiple formats and folders across a network, and provides a single, user-friendly interface to do everything at once or in scheduled steps

INTERSYSTEMS
www.intersystems.com

IRIS—a unified solution that provides a comprehensive and consistent set of capabilities spanning data management, interoperability, transaction processing, and analytics—eliminating the need to integrate multiple development technologies so applications require less code, fewer system resources, and less maintenance

KORE TECHNOLOGIES
www.koretech.com

Kourier Integrator—Kore's flagship enterprise integration and data management suite, providing extract, transform, and load and enterprise application integration capabilities, helps organizations extend the value and functionality of enterprise applications by integrating and connecting with disparate databases and best-in-class applications

LICENSEFORTRESS
www.licensefortress.com

LicenseFortress—database software license management and audit protection services comprised of four levels (Discovery, Standard, Premier, and Legal) to help organizations enhance their understanding of software licensing and pricing practices, reduce risk, and "eliminate the surprise factor"

LOOKER DATA SCIENCES
https://looker.com

Looker Platform—designed to flex and scale as needs for information change and expand, the modern platform for data continues to build upon its BI application and adds stronger enterprise-class features, enhanced tools, and powerful department-specific plug-and-play applications to help users access, analyze, and take action

MAPR TECHNOLOGIES
www.mapr.com

MapR Converged Data Platform—a platform for artificial intelligence and analytics that enables enterprises to inject analytics into their business processes to increase revenue, reduce costs, and mitigate risks, while addressing the data complexities of high-scale and mission-critical distributed processing

MARIADB
https://mariadb.com

MariaDB TX—an open source database fueled by community innovation that offers Oracle Database compatibility, sharding, temporal tables, and point-in-time rollback, and uses purpose-built storage engines to maximize the storage efficiency and query performance of applications with traditional and/or non-traditional workloads

MARKLOGIC
www.marklogic.com

MarkLogic—an operational and transactional enterprise NoSQL database platform designed for speed and scale, recently introduced the MarkLogic Query Service, a new way to give customers elasticity in the cloud

MELISSA
www.melissa.com/data-integration/talend

Data Quality Components for Talend—offers built-in solutions for Talend Opera Studio for Data Integration, including Personator, Melissa's all-in-one ID verification, data completion, and data-enrichment tool, along with Global Address Verification to clean, standardize, and transliterate addresses in more than 240 countries

MEMSQL
www.memsql.com

MemSQL—offers a single database that can easily support real-time decision making and insight-driven customer experiences combining massively scalable ingest, the ability to scan billions of rows per second across thousands of users, and a familiar relational structure for user- and developer-friendly access

MICROSOFT
https://azure.microsoft.com

Azure—a set of cloud services that developers and IT professionals use to build, deploy, and manage applications through a global network of data centers, with additional support provided by integrated tools, DevOps, and a marketplace for building anything from simple mobile apps to internet-scale solutions

MICROSTRATEGY
www.microstrategy.com

MicroStrategy—an enterprise analytics platform that offers out-of-the-box visualizations, intelligent recommendations for content, prompts for dossiers, a native MicroStrategy Library app for smartphones, and a new enhanced mapping for conducting geospatial analytics with Mapbox, a leading location data platform for mobile and web applications

MONGODB
www.mongodb.com

MongoDB Atlas—engineered by the same team that builds the MongoDB NoSQL database platform, the MongoDB Atlas database-as-a-service incorporates best practices developed from real-world use cases at thousands of customer deployments from startups to the Fortune 100

NUMERIFY
www.numerify.com

Numerify—contains a suite of specialized IT applications powered by an underlying analytical platform with each analytical application focusing on a specific business area of IT—including service, asset, and project management

NUXEO
www.nuxeo.com

Nuxeo Platform—a modern content services platform that manages both traditional content and rich media assets to unlock the full value of all of a company's digital content—anywhere within the company

OPENTEXT
www.opentext.com

OpenText Legal Center—offers a process-centric approach to the legal field for client onboarding, external sharing and collaboration, and document management, and is designed to leverage and extend existing DM repositories such as OpenText eDOCS, whether on-premise or in the cloud

ORACLE
www.oracle.com

Oracle Autonomous Database Cloud—powered by Oracle Database 18c, the next generation of the industry-leading database, Oracle Autonomous Database Cloud offers total automation based on machine learning and eliminates human labor, human error, and manual tuning

PAXATA
www.paxata.com

Adaptive Information—a big data prep platform aimed at democratizing the data preparation process, offering new features such as a Adaptive Workload Management capability, which delivers an elastic resource allocation service on a number of orchestration frameworks, including Microsoft Azure HDInsight, Kubernetes, and Apache Hadoop YARN

PERCONA
www.percona.com

Percona Monitoring and Management—combines best-of-breed tools in a single, virtual appliance, along with Percona-developed query analytics, administration, API, agent, and exporter components, to monitor and provide performance data for MySQL and MongoDB variants

PREMIUMSOFT CYBERTECH
www.navicat.com

Navicat Premium—a multi-connection database development tool that allows users to connect up to six databases within a single application, including MySQL, MariaDB, SQL Server, SQLite, Oracle, and PostgreSQL, to create access to all databases at once

PROGRESS
www.progress.com

Progress OpenEdge—an application development platform that helps simplify the delivery of mission-critical business applications with a focus on increasing the ability to always be on, fortify applications through enhanced security, and keep accurate data flowing through the organization

PURE STORAGE
www.purestorage.com

Pure Storage FlashBlade—a petabyte-scale storage system for unstructured and operational data that is designed with four key innovations to overcome the limitations of conventional clustered NAS systems: high-performance storage devices, unified network, distributed storage operating system, and administrative simplicity

PYTHIAN
https://pythian.com

Pythian Kick Analytics-as-a-Service—a cloud-native analytics platform that solves the data silo problem, taking the hard work out of capturing, curating, and preparing data for consumption by a range of users and systems

QLIKTECH
www.qlik.com

QlikSense—lets users discover, search, and explore across all their data, pivoting their analysis when new ideas surface, and provides flexibility with a cloud-ready data analytics platform that supports the full spectrum of BI use cases

QUBOLE
www.qubole.com

Qubole Data Service—provides a single platform for ETL, reporting, ad hoc analysis, stream processing, and machine learning to help data teams be more productive and reduce the costs of their data initiatives while taking full advantage of the elasticity and scale of the cloud

QUEST SOFTWARE
www.quest.com

SharePlex—a multi-platform replication and real-time data integration solution that supports organizations' increasingly heterogeneous database environment used by Fortune 500 for migrations and upgrades, load balancing, supporting BI and analytics, and much more

RED HAT *(acquired by IBM)*
www.redhat.com

Ansible Automation—simple, agentless IT automation technology that can improve business processes, migrate applications for better optimization, and provide a single language for DevOps practices across the organization using a simple, human-readable language that anyone in IT can understand

REDIS LABS
https://redislabs.com

Redis Enterprise—combines the advantages of world-class database technology and the innovation of a vibrant open source community to provide high availability in the form of Active-Active and Active-Passive geographically distributed architectures, with high performance and built-in search capabilities

REDPOINT GLOBAL

www.redpointglobal.com

RedPoint Customer Engagement Hub—delivers a unified view of each customer, in-line analytics to determine next-best actions, and intelligent orchestration to personalize engagement across the enterprise, and recently added new master data management capabilities

REVELATION SOFTWARE

www.revelation.com

OpenInsight 10—Revelation's flagship product is a NoSQL database development suite that provides Windows, Web 2.0, and .NET tools to develop and deploy mission-critical applications along with a redesigned integrated development environment, 64-bit architecture integration, new security updates, and more

ROBIN SYSTEMS

https://robin.io

ROBIN Hyper Converged Kubernetes Platform—hyper-converged Kubernetes is a software-defined application orchestration framework that combines containerized storage, networking, compute (Kubernetes), and the application management layer into a single system, allowing users to do self-service deployment of big data, databases and AI/ML, share entire experiments among team members, and more

ROCKET SOFTWARE

www.rocketsoftware.com

UniVerse—a component of the MultiValue application platform, is a fast, flexible data server for developing enterprise app, and its variable length, table-within-table architecture means speedy data access and low maintenance with user interfaces for Windows, Linux, and UNIX

SAP

www.sap.com

HANA—combines an ACID-compliant database with application services, high-speed analytics, and data acquisition tools in an in-memory platform with built-in application services that support the development and deployment of new business apps to deliver insight from big data and IoT

SAS INSTITUTE

www.sas.com

Viya—complements SAS, augmenting the SAS Platform to enable everyone—data scientists, business analysts, developers, and executives alike—to collaborate and achieve innovative results faster to derive new insights like never before

SOFTWARE DIVERSIFIED SERVICES (SDS)

www.sdsusa.com

IronSphere—an ISCM solution for the IBM mainframe that provides complete visibility of system risks, IronSphere continuously monitors the mainframe, automatically performing security scans and looking for system vulnerabilities, insufficient settings, and modified operands

SEARCH & CONTENT ANALYTICS *(part of Accenture Applied Intelligence)*

www.searchtechnologies.com

Data Lake Solutions—leverages the company's expertise in Hadoop, Cloudera CDH, Cassandra, Elastic Stack, Solr, Hortonworks, Microsoft Azure, Amazon Web Services, and other big data platforms to offer a wide range of flexibility and security options for a data lake

SENTRYONE

www.sentryone.com

SQL Sentry for SQL Server—allows organizations to collect and present the most actionable performance metrics and alerts, see important events, and cross-reference them using an Outlook-style calendar to resolve issues by doing detailed analysis in the same tool used for monitoring and alerting

SINEQUA

www.sinequa.com

Cognitive Search & Analytics Platform—offers broad connectivity to all enterprise data sources and combines disruptive analytical tools, including natural language processing and machine learning, to increase the enterprise IQ and continuously improve its ROI

SISENSE

www.sisense.com

Sisense—simplifies business analytics for complex data by providing a complete solution for preparing, analyzing, and visualizing big or disparate datasets with agility that allows business users with no technical background to get the accurate intelligence at the very moment it is needed

SNAPLOGIC

www.snaplogic.com

SnapLogic Enterprise Integration Cloud—accelerates digital transformation to connect the application and data landscape, and empowers business users and IT—with intuitive, scalable, and connected integration along with new features such as integration with GitHub and support for Mesosphere to automate critical elements of continuous integration and delivery

SNOWFLAKE

www.snowflake.com

Snowflake Data Warehouse—a new SQL data warehouse from the ground up for the cloud designed with a patented new architecture to handle all aspects of data and analytics to deliver the performance, simplicity, concurrency, and affordability not possible with other data warehouses

SOLARWINDS

www.solarwinds.com

Database Performance Analyzer—monitors the entire environment on premises, in the cloud, or as a service along with helping to identify performance issues and eliminate database bottlenecks using multi-dimensional views to answer the who, what, when, where, and why

SOUTHBANK SOFTWARE

www.dbkoda.com

dbKoda—an open source integrated development environment for MongoDB includes a performance lab that offers a graphical real-time view of MongoDB server performance, allowing users to drill into poorly performing commands, and an index advisor to help users optimize those commands

SPLUNK

www.splunk.com

Splunk Enterprise—a platform that enables operational intelligence from machine-generated data and provides a range of search, visualization, and prepackaged content cases to support organizations in discovering and sharing insights easier and faster

SQREAM

https://sqream.com

SQream DB—a fully-featured GPU-accelerated data warehouse, capable of handling the most complex queries featuring comprehensive ANSI SQL support to load, store, and analyze data by using advanced columnar techniques, along with high-throughput processors

STRIIM

www.striim.com

Striim—a patented, enterprise-grade platform that offers continuous real-time data ingestion, high-speed in-flight stream processing, and sub-second delivery of data to cloud and on-premise endpoints, continuously delivering data that can be immediately available to high-value operational workloads

SYNCSORT

www.syncsort.com

DMX-h—designed to help users achieve their modern data strategy objectives, DMX-h features a single interface for accessing and integrating enterprise data sources—batch and streaming—across Hadoop, Spark, Linux, UNIX, or Windows, whether on premises or in the cloud

TABLEAU SOFTWARE

www.tableau.com

Tableau—available as desktop, server, and hosted software, the signature platform allows users to connect, explore, and visualize their data to achieve insights through automated table and join recommendations powered by ML algorithms that simplify the search for the right data for analysis

TALEND

www.talend.com

Talend—simplifies and automates big data integration with graphical tools and wizards that generate native code, enabling teams to start working with Apache Hadoop, Apache Spark, Spark Streaming, and NoSQL databases, for cloud or on-premises today

TERADATA
www.teradata.com

Teradata Vantage—enables enterprises to uncover actionable answers to the toughest business questions by tightly integrating the best analytic functions and engines to provide a scalable, agile platform that helps them to drive business value

TIGERGRAPH
www.tigergraph.com

TigerGraph—a fast graph analytics platform for the enterprise offers new features that include seamless integration with popular databases and storage systems, support for Docker and Kubernetes containers, availability on the Amazon Web Services Marketplace and Microsoft Azure, and a new graph algorithm library

TRIFACTA
www.trifacta.com

Trifacta Wrangler—accelerating data preparation for analytics and machine learning with no coding required, Wrangler makes data cleaning, blending, and structuring more intuitive so that even the most challenging Excel, CSV, or JSON files can be wrangled in minutes

VERTICA
www.vertica.com

Vertica Analytics Platform—designed and built on a tested, reliable distributed architecture and columnar compression to deliver blazingly fast speed for use in data warehouses and other big data workloads where speed, scalability, simplicity, and openness are crucial to the success of analytics

VMWARE
www.vmware.com

VMware Cloud Foundation—provides integrated cloud infrastructure (compute, storage, networking, and security) and cloud management services to run enterprise applications in both private and public environments by leveraging a common infrastructure and consistent operational model

VOLTDB
www.voltdb.com

VoltDB—an in-memory, scale-out operational database, built to help organizations create high-velocity applications, that offers increased support for the Hadoop ecosystem, expanded SQL support, and a new management center without compromising on ACID requirements

WHERESCAPE
www.wherescape.com

Wherescape Automation—on-premise, in the cloud, or a combination of both, WhereScape data automation solutions provide out-of-the-box best practices, optimized native code, and features for popular target data platforms, including Amazon Redshift, Microsoft SQL Server, Microsoft Azure, Oracle, Snowflake, and Teradata

YELLOWFIN
www.yellowfinbi.com

Yellowfin—a BI and analytics platform aimed at solving real enterprise analytics challenges and helping business people understand not only what happened, but why through a series of solutions offering data discovery, data prep, and more

ZALONI
www.zaloni.com

Zaloni Data Lake Management Platform—a comprehensive, integrated solution that operationalizes data along the entire pipeline, providing automation of repeatable management tasks and processes and the ability to centrally manage all enterprise data sources regardless of location

ZOOMDATA
www.zoomdata.com

Hadoop Big Data Visualization—a Hadoop data visualization tool that can connect directly to HDFS as well as to SQL-on-Hadoop technologies, such as Impala, Hive, SparkSQL, and Presto, and connect to big data sources such as search, streaming, and NoSQL via smart connectors

PRODUCT SPOTLIGHT

Aerospike

Brian Bulkowski, Founder and CTO

Aerospike is trusted by leading enterprises around the world to help them confidently deploy mission-critical, strategic operational applications that make digital transformation possible. Aerospike's vision is to make it easy and affordable for companies of all sizes to build next-generation data systems like those built internally by the largest internet-scale companies like Google and Facebook.

Our enterprise-grade database is deployable anywhere, delivers unmatched uptime, predictable performance, and exceptionally low TCO. Aerospike has customer deployments that have run for years with no service disruption, handling hundreds of terabytes of data, supporting trillions of transactions per month, with sub-millisecond latency.

Aerospike is used in Financial Services, Banking, Telecommunications, Technology, Retail, E-commerce, Adtech, Martech and Online Gaming. Powered by a patented Hybrid Memory Architecture™ and autonomic cluster management, the Aerospike database is ideal for Fraud Prevention, Digital Wallet, Online Brokerage, Telco Charging & Billing, Messaging & Chat, Recommendation Engines, Real-Time Bidding, and other applications that require the highest possible uptime, performance and scale.

Aerospike customers include Adobe, Airtel, FlipKart, Kayak, Nielsen, Nokia, and Snap.

Aerospike
www.aerospike.com

BackOffice Associates

Rex Ahlstrom, Chief Strategy and Technology Officer

On behalf of the BackOffice Associates team worldwide, it is an honor to accept DBTA's recognition of our Information Governance Cloud solution for its Trend-Setting Products for 2019. As many of today's leading global enterprises undergo digital transformations, they are also embarking on a data transformation as they shift onto cloud-based business suites and must manage more structured and unstructured data than ever before. As the pace and complexity of doing business is accelerating exponentially, having reliable, business-ready data available across the full enterprise system landscape is critical to remain competitive.

With these evolving market dynamics in mind, BackOffice Associates developed Information Governance Cloud to help organizations to get more value from their vast amounts of data by validating its accuracy and relevancy, as well as using it to drive key business outcomes. The solution crowdsources contributions and knowledge from business and data experts with guidance from artificial intelligence and machine learning to yield trusted data.

Empowering business and IT stakeholders alike, Information Governance Cloud offers the ability to comprehensively set and enforce data policies across an enterprise through smart automation. By using Information Governance Cloud, organizations can define and streamline operational policies and business rules for data, align them with the organization's overall business strategies, and oversee policy enforcement across all systems including BackOffice Associates' Data Stewardship Platform™.

As the industry continues to expand, BackOffice Associates remains committed to delivering innovative solutions like Information Governance Cloud, which harness the power of information governance, allowing business users to use their data to execute on mission-critical business imperatives.

BackOffice Associates
www.boaweb.com

PRODUCT SPOTLIGHT

TREND-SETTING PRODUCTS for 2019

Bradmark Technologies, Inc.

SURVEILLANCE DB™ PERFORMANCE MONITORING & EVENT MANAGEMENT SOLUTION

Edward Stangler
R&D, Product Manager, Surveillance

Enterprises are mixing cloud and on-premises infrastructure in ways that make sense for them. Performance monitoring solutions need to handle this mix of infrastructure, as well as the distributed nature (technical and logical) of many businesses across geographic zones.

ESSENTIAL VISIBILITY ACROSS YOUR GROWING IT INFRASTRUCTURE MIX

Bradmark's Surveillance DB™ for enterprise environments provides the most comprehensive, real-time monitoring and proactive event management solution available. With support for SAP-Sybase, MS-SQL Server, Oracle, Informix and DB2 databases running on UNIX, Linux and Windows servers, Surveillance DB supports your growing IT infrastructure by providing:

Seamless integration between cloud-accessible monitoring and on-premises (or cloud) databases. Quickly provide secure global web access to performance monitoring resources in data centers, whether it's in-house or remote locations.

Access to local, distributed recent history across many sites. Utilizing three-tier, distributed metric data, Surveillance DB provides quick local access to recent history, without using WAN bandwidth or having contention on centrally-located resources. So, for industries with multiple store locations or business organizational units—many key metrics are available from the recent past using data that is stored locally at each site, rather than having to query a central database repository over the network.

And to ensure that global infrastructures are managed efficiently and effectively, our future roadmap highlights alerts augmented with machine learning and reactive feedback—as good alerts turn into important details, and important details turn into faster solutions for problems.

Bradmark Technologies, Inc.
www.bradmark.com/surveillance

Cambridge Semantics ANZOGRAPH

Steve Sarsfield,
VP of Product

Today's most challenging analytics tasks center around the connections in the data. If you're confronted with a difficult process when trying to create corporate knowledge graphs, understand buyer intent and influence, detect fraud or build a recommendation engine, a data warehouse style analytics solution that focuses on this type of connected analytics is imperative to your success. You need a graph database like AnzoGraph™ from Cambridge Semantics.

AnzoGraph is a native, Massively Parallel Processing (MPP) distributed Graph OLAP (GOLAP) database, providing fast advanced analytics at big data scale. Data analysts, enterprise architects and application developers can use AnzoGraph to build and execute data warehouse analytics, graph analytics and inferencing, all in one award-winning database. Graph OLAP databases offer better performance on deep, long-running queries when compared to graph OLTP databases like Neo4j, AWS Neptune and other OLTP systems.

AnzoGraph offers graph analytic functionality to enable new levels of data-driven insight to help drive new business opportunities, minimize costs and increase competitiveness. AnzoGraph is capable of loading, updating, querying and analyzing huge amounts of data at speeds up to 100x faster than other graph databases, and our benchmarks prove it.

AnzoGraph supports both labeled property graphs (using W3C RDF* proposed standard) and semantic graphs with advanced analytics functionality, including graph views, named queries, analytic/data mining/reporting functions, inferencing and conditional expressions. You can perform special graph functions like pagerank, shortest path, all paths and many more. It's available now for a free download and trial.

AnzoGraph by Cambridge Semantics
www.anzograph.com

PRODUCT SPOTLIGHT

Datavail

Eric Russo, SVP, Database Services

Datavail Delta Monitoring Expands to MySQL and Oracle on Windows Server.

As the core component of data-driven software, your databases need to be continuously managed to make your business more agile and less susceptible to risk. This is why Datavail developed our custom-built software, Delta, a world-class database monitoring solution.

We're excited to share that our platform has expanded to provide proactive monitoring for MySQL and Oracle. Now, DBAs managing databases across cloud and on-premise platforms can set aside tedious, manual database monitoring, and rely on Delta to provide alerts for the most pressing issues.

Delta can report on your most critical areas:

- Availability
- Configuration
- Backups
- Scheduled Jobs
- Security and Compliance
- Performance
- Capacity Planning

With intelligent alerts that dispatch the right level of response, Datavail Delta ensures your most pressing issues receive immediate attention from database administration professionals.

Find out why 225+ Datavail customers are using Delta to monitor more than 5,000 servers and 125,000 databases. To learn more about Delta, visit our website, or contact Datavail today at 866-239-9538.

Datavail
www.datavail.com

Datawatch Angoss

Michael Rowley, Director Global Product Marketing

Datawatch Angoss is a global leader in delivering advanced analytics to businesses looking to monetize data with a focus on logistics, risk, marketing and sales. Datawatch Angoss removes the complexity inherent to predictive analytics and machine learning by offering a platform that is both intuitive to use and rich in features. These features include intuitive workflows with drag-and-drop functionality that train, test and deploy predictive models for quicker time-to-insight in model outputs.

KnowledgeSTUDIO for Apache Spark provides a productivity tool for Spark users, allowing them to interact with Spark via a graphical user interface to generate error-free code that may be used in production scripts. Data science teams that are modeling in a Big Data environment and outside of it can use Angoss KnowledgeSTUDIO for Apache Spark to efficiently build analytic workflows using datasets of all sizes. It effectively leverages Apache Spark's ability to operate on datasets with extremely large numbers of records, while improving SparkSQL and SparkML queries on datasets that have thousands of columns. Very small datasets can also be efficiently included in Spark analytical workflows. Support for Hive, Parquet, and distributed CSV is also featured.

Leading global organizations in financial services, insurance, retail and technology rely on Datawatch Angoss to grow revenue, increase sales productivity and improve marketing effectiveness while reducing risk and cost. Datawatch Angoss solutions include scenario comparison, cloud-based model deployment and real-time scoring, and sentiment data analytics to capture trending conversations about consumer brands and product.

Datawatch Angoss
https://www.datawatch.com

PRODUCT SPOTLIGHT

TREND-SETTING PRODUCTS for 2019

Delphix

Eric Schrock, CTO

Today's modern businesses need efficient, secure access to data to innovate faster to keep pace with consumers demands and avoid disruption.

An overwhelming majority of Fortune 1000 executives fear disruption from nimble, data-driven competitors, and say the inability to be agile and compete on data is the most pressing challenge facing their business today. Their fear is not unwarranted. Over the last decade, nearly three-quarters of Fortune 1000 companies have been replaced, despite aggressive investments in software to fend-off more agile, tech-savvy competitors.

Access to data for business is no longer a forward-thinking trend, it's a prerequisite to survival in today's fast-moving market. However, data "friction" can keep companies from moving fast while prioritizing security. Just one data breach can cause profound damage to corporate reputations and customer perceptions. Businesses are under tremendous pressure to deliver data for innovation while protecting it from crippling security breaches.

At Delphix, our mission is to solve these pain points by eliminating data friction. As the critical foundation for any strong DataOps strategy, we provide automated, secure, self-service access to data for the people who need it to drive innovation and business agility. The Delphix Dynamic Data Platform virtualizes and secures data to make it lightweight and accessible for the people and teams who need it, regardless of their size.

Hundreds of the largest and most valuable brands in the world depend on Delphix to accelerate the time to market for their key initiatives, and greatly enhance the quality of their offerings, including one-third of Fortune 100 companies.

Delphix
www.delphix.com

Denodo Technologies

Ravi Shankar, CMO Data Virtualization

Organizations are moving data to the cloud, driven by efficiencies of scale; reductions in the total cost of hardware, software, and the entire IT infrastructure; flexible, usage-based payment structures; and the ability to provision new IT capabilities in minutes rather than months.

Migrating data from on-premises systems requires a certain amount of downtime, however, and many organizations are discovering that their employees', customers' and partners' tolerance for downtime is rapidly approaching zero. Such organizations need a multilocation architecture, which is a data architecture that enables access to both on-premises and cloud systems simultaneously. With a multilocation architecture, migrations can take place without downtime, and without users' knowledge that a migration has taken place.

Multilocation architectures enable location transparency and data abstraction, and as such they support the world's most advanced data infrastructures. Beyond hybrid on-premises/cloud environments, multilocation architectures also support environments that extend across multiple cloud providers. Such environments may leverage AWS for some applications, Azure for others, and Google Cloud for still others, including emerging AI applications.

The Denodo Platform supports multilocation architecture through data virtualization, which enables real-time data integration between on-premises systems and cloud systems, or between different cloud systems, all without moving the data. Data virtualization establishes a universal data access layer between the sources. To access the data, no matter where it resides, a user or consuming application simply needs to access the data virtualization layer, which gets a view of the requested data in real time. The Denodo Platform provides the most advanced data virtualization solution, enabling all the benefits of a modern multilocation architecture!

Denodo Technologies
www.denodo.com

PRODUCT SPOTLIGHT

Empolis

Stefan Wess, CEO

With the cloud-based solution Empolis Service Express, technical customer service is successfully digitized immediately, with no large investments. Empolis Service Express provides global, scalable access to distributed service information for service personnel, partners, and customers. Empolis Service Express uniquely combines intelligent search and guided troubleshooting based on artificial intelligence.

Magical service moments are created with Empolis Service Express, therefore allowing service staff to deliver excellent customer service and become real service heroes. They have access to all relevant service information with a single action, distributed across a wide range of tools. They find the right information immediately, even in large documents. All employees share their experiences quickly and easily with their colleagues and always have the full range of information available online and offline.

Integrated analytics tools instantly make it clear where problems exist, how customers can be assisted and taps into the potential for improvement. Thus, information quality is continuously enhanced.

Empolis Service Express is also available as a mobile service app for smartphone or tablet assistance for support field technicians as they conduct their service appointments. Since the app has full offline capability, all relevant information and the company's entire service knowledge can be accessed on site, at any time, without an internet connection. All of this, in an instant, free of tedious data input, simply by selecting options or via voice commands, supported by comprehensive ergonomic design and user guidance. The Empolis Service Express mobile app is available in the app store.

Empolis
www.empolis.com

erwin EDGE

Mariann McDonagh, Chief Marketing Officer

We developed the erwin EDGE platform to deliver an "enterprise data governance experience," so the modern business can accelerate the transformation of mission-critical data into accurate and actionable insights.

The erwin EDGE ensures collaboration between IT and the business to discover, understand and unlock the value of data both at rest and in motion. With data governance at the hub, it brings together business process, enterprise architecture, data mapping and data modeling to simplify the complete data management and governance lifecycle. And it features the broadest set of metadata connectors and automated code generation, data mapping and cataloging tools available today.

This single, integrated solution makes it possible to gather business intelligence, conduct IT audits, ensure regulatory compliance and accomplish any other organizational objective by fueling an automated, high-quality and real-time data pipeline.

With the erwin EDGE, organizations can:

- Discover data: Identify and integrate metadata from various data management silos.
- Harvest data: Automate the collection of metadata from various data management silos and consolidate it into a single source.
- Structure data: Connect physical metadata to specific business terms and definitions and reusable design standards.
- Analyze data: Understand how data relates to the business and what attributes it has.
- Map data flows: Identify where to integrate data and track how it moves and transforms.
- Govern data: Develop a governance model to manage standards and policies and set best practices.
- Socialize data: Enable stakeholders to see data in one place and in the context of their roles.

erwin, Inc.
www.erwin.com

PRODUCT SPOTLIGHT

TREND-SETTING PRODUCTS for 2019

Franz Inc. ALLEGROGRAPH—GRAPH DATABASE FOR AI KNOWLEDGE GRAPHS

Jans Aasman, CEO

Artificial intelligence (AI) is one of the top investment areas for companies looking to improve ROI on operations and products, and to create customer 360 views. Using AI to create "Enterprise Knowledge" and link it across the Enterprise to create a "Knowledge Graph" is a key differentiator for companies in an ever-increasing competitive landscape. The foundation for Knowledge Graphs and Artificial Intelligence lies in the facets of semantic technology provided by Franz's AllegroGraph database. Semantic Graph databases, such as AllegroGraph, provide the core technology environment to enrich and contextualize the understanding of data. The ability to rapidly integrate new knowledge is the crux of the Knowledge Graph and depends entirely on semantic technologies.

An early innovator in Artificial Intelligence, Franz Inc. is a leading supplier of Knowledge Graph solutions with Semantic Graph Database technology as the foundation. If you really want to develop your corporate Knowledge Graph and address complex Artificial Intelligence problems, you need a data system that goes beyond just data. You have to create a system that can link to anything outside your own predefined parameters—and that can learn from previous experiences. That is where a Semantic Graph Database, like AllegroGraph, comes into the picture.

Franz Inc. provides a variety of services as part of its Knowledge Graph platform solution: from architectural consulting and technical seminars to training. Franz's flagship product, AllegroGraph, provides the necessary power and flexibility to address high-security data environments such as HIPAA access controls, privacy rules for banks, and security models for policing, intelligence, and government.

Contact Franz Inc. to unleash the potential of your Company's Knowledge Graph.

Franz Inc.
https://franz.com/

Griddable

Robin Purohit, CEO

As enterprises embrace hybrid cloud as the new long-term reality, database architects need a simple approach to migrate enterprise data to the cloud, re-engineer for cloud-first architectures, and ensure portability of data across multiple public clouds to avoid lock-in.

Griddable migrates and modernizes legacy data seamlessly across heterogeneous clouds and database types. The Griddable platform is a cloud-first product, providing data synchronization with scale up, scale out, high availability features out of the box. Most importantly, Griddable provides an automated user experience that requires nothing more than connecting the source and target databases to the grid. Once connected, Griddable automates schema copying, initial load, and subsequent change data capture in a continuous process that requires no user intervention.

Many use cases today require modernizing data to multiple targets simultaneously. Whether replacing a legacy Oracle database or re-architecting for microservices, the Griddable data pipeline efficiently synchronizes to multiple destinations with low latency while supporting data customizations at each destination.

The Griddable policy engine provides user-friendly controls for modernizing and transforming data in transit. Griddable masks or encrypts any number of individual data elements using separate masking algorithms or encryption keys. It also filters and replaces data values, or selectively removes entire rows or columns, with an easily-defined user policy. Using Griddable.io's selective filtering and synchronization, cloud operators can accelerate movement of key data to best-fit platforms. Decisions regarding data placement can be made dynamically during production and adjusted to operational parameters.

To learn more about Griddable.io, download our white paper https://griddable.io/hybrid-cloud/.

Griddable
griddable.io

PRODUCT SPOTLIGHT

InfluxData

Mark Herring, CMO

Time series has been the fastest growing database category for the last two years, according to DB-Engines. This growth is being fueled by two major industry trends—the rapid instrumentation of the physical world, driven by increasing investment in IoT systems, and an explosion in the software world of cloud-native applications and services, all of which are being instrumented for real-time visibility and control. This "Age of Instrumentation" has generated major demand for purpose-built time series platforms that can support the critical requirement for real-time processing of the myriad metrics and events that deliver insights and competitive advantage to data-driven organizations.

InfluxData provides the leading time series platform, built from the ground up, for analyzing metrics and events for DevOps and IoT applications. Whether the data comes from humans, sensors, or machines, InfluxData empowers developers to build next-generation monitoring, analytics, and IoT applications faster, easier, and to scale to deliver real business value, quickly. Based in San Francisco, InfluxData's more than 450 customers include Cisco, eBay, IBM and Siemens.

InfluxData
www.influxdata.com

InterSystems IRIS DATA PLATFORM: INTUITIVE, RELIABLE, INTEROPERABLE, AND SCALABLE

Carlos Kuhl Nogueira, General Manager, Data Platform Initiatives

Businesses need new applications to meet the growing demands of society's advancing technology—applications that are smarter, faster, and that can scale more quickly and cost-effectively. The solution? InterSystems IRIS Data Platform™, a single, unified platform that provides all of the following capabilities to efficiently accommodate increasing workloads and data sizes:

Data Management: InterSystems IRIS is an ultra-high performance, horizontally scalable, multi-model database. It stores and accesses data modeled as objects, schema-free data, relational data, and multi-dimensional arrays in a single, highly efficient representation. It simultaneously processes both transactional and analytic workloads in a single database at very high scale, eliminating latencies between event, insight, and action.

Interoperability: A comprehensive, integration platform provides application integration, data coordination, business process orchestration, composite application development, API management, and real-time monitoring and alerting capabilities to support the full spectrum of integration scenarios and requirements.

Data Analytics: A powerful, open analytics platform supports a wide range of analytics and is able to analyze real-time and batch data simultaneously at scale. Developers can embed analytic processing into business processes and transactional applications, enabling sophisticated programmatic decisions based on real-time analyses. Natural language processing capabilities allow developers to extract meaning and sentiment from unstructured text.

Cloud Deployment: Automated "cloud-first" deployment options simplify public cloud, private cloud, on premise, and virtual machine deployments and updates.

InterSystems IRIS redefines high performance for application developers, systems integrators and end-user organizations who develop and deploy data-rich and mission-critical solutions.

InterSystems
InterSystems.com/IRIS

PRODUCT SPOTLIGHT

TREND-SETTING PRODUCTS for 2019

Kore Technologies

Mark Dobransky,
Co-Founder and
Managing Partner

Kourer Integrator is Kore's flagship product for advanced Enterprise Application Integration (EAI) and Extract, Transform and Load (ETL) solutions.

Kore clients, partners and VARs have been using Kourier for years to extend the value and functionality of their enterprise applications by creating near real-time data warehouses from disparate data sources and through asynchronous integration with best-in-class applications.

However, the speed of business today often requires more synchronous, real-time integration between applications. RESTful Web Services are now the technology of choice and Kourier is evolving to meet these demands.

Our latest release of Kourier continues to improve on its ability to create REST APIs that give third-party applications secure, rated access to your MultiValue data (REST as server). Plus, now you can integrate to other third-party applications via their REST endpoints (REST as a client). Using Kourier's new application "Connectors" you can use REST to talk to other applications from your MultiValue system (e.g., CRM, Service Management and eCommerce) just as easily as they can talk to your system.

Integration can be hard, so we endeavor to make it easier using our "clicks not code" approach. Developers are more productive creating REST-based integrations using Kourier's template-based APIs and application connectors.

Kore understands the challenges faced by developers and technology is constantly changing. We strive to provide efficient and easy-to-use software and to react quickly to the needs of our partners and clients so that we continue be a trend-setting solution provider.

Kore Technologies
www.koretech.com
Or call: 866-763-5673

Melissa TIRED OF THE ETL GRIND? WORK FASTER AND SMARTER WITH MELISSA AND TALEND

Bud Walker,
VP of Enterprise
Sales & Strategy

In the era of big data, DBAs spend much of their time transferring, migrating, and cleaning disparate data. Melissa makes this more efficient by adding data quality into the workflow to help DBAs work faster and smarter with clean data. We've teamed up with Talend, a leading data integrations firm, to offer a faster, easier way to connect and clean data with over 900 built-in connectors, advanced string manipulations, and a unified metadata repository.

Through the Talend Open Studio platform, Melissa offers built-in data quality components including Personator, our ID verification tool. Personator matches name to address, email, and phone number, and verifies national ID and age, in real-time. The tool also enriches your contact data with demographics such as consumer or business name, address, phone or email, gender, household income, FIPS Code, deceased info, and much more.

We also offer our Global Address Verification component to clean, standardize, transliterate, and verify deliverability of addresses in over 240 countries.

The combined solution brings agility in the data pipeline as data processing and data quality controls can be applied inflight, on the cloud, or on-premise in real-time or batch.

Boost your productivity and easily clean and integrate all your data with the combined power of Melissa's data quality solutions and Talend's robust open source ETL tool. Free trials, unlimited tech support, and a 120-day ROI guarantee are available.

Melissa
www.melissa.com/dbta-etl

PRODUCT SPOTLIGHT

Navicat

Navicat

Navicat is an industry-leading and award-winning database management and development solution. Our first product—Navicat for MySQL—was launched in 2002. To help our users stay competitive in today's business world, we continue to improve our products and add new database support. Today, our top-rated prod-uct, Navicat Premium, supports 7 databases within a single application, including MySQL, MariaDB, MongoDB, SQL Server, SQLite,Oracle and PostgreSQL, and supports 12 languages.

Navicat Monitor is our new product line. "We know how hard and important it is to monitor your database performance," said Ken Lin, CEO at PremiumSoft. "With Navicat Monitor, we want our users to be able to keep track of how their database is used and alerted if there are any threshold breaches so as to ensure their database performs to the highest standards. Now, Navicat Monitor supports MySQL and MariaDB databases. We will continue to add new database support and new features in the future."

With more than 16 years of providing database management solutions and more than 40% of Fortune 500 companies counting on Navicat every day, we will continue to offer world-class customer support and exciting new features. We are always working on new features and innovations to help your business gain a competitive edge.

Download a free 14-day free trial of :

Navicat Premium:
https://www.navicat.com/download/navicat-premium

Navicat for MongoDB:
https://www.navicat.com/en/download/navicat-for-mongodb

Navicat Monitor:
https://www.navicat.com/en/download/navicat-monitor

Navicat
www.navicat.com
Navicat is the choice of over 3 million database users all around the world. Over 160,000 registered customers across 7 continents and 138 countries have chosen our products.

Percona

Michael Coburn, Product Manager for Percona Monitoring and Management

Percona Monitoring and Management (**PMM**) is a free, open source database monitoring and performance optimi- zation tool for MySQL, MongoDB, and PostgreSQL. PMM builds on other open source projects including Grafana, Prometheus, and Consul and is deployed in your environment on your equipment for maximum security and reliability. PMM gives you two views on your database performance: Query Analytics provides information about queries running in your database, and Metrics Monitor plots database metrics over time.

Query Analytics helps you optimize database performance by identifying those queries in MySQL and MongoDB that consume the most amount of resources, providing clarity on which queries deserve attention first. Query Analytics highlights the following query attributes:

- Query response time, locks, rows sent, rows examined
- Percona Server for MySQL only: InnoDB operations: reads, waits, Temporary Table usage
- MySQL query plan information via EXPLAIN in table and JSON format so you can evaluate the level of optimization.

Metrics Monitor displays database activity over time, including:

- MySQL, MongoDB, and PostgreSQL: Queries Per Second (QPS), replication, storage engine-specific (InnoDB, WiredTiger, MyRocks, RocksDB, MMAPv1, MyISAM, In-Memory, TokuDB, and Aria engines)
- System-level resources: CPU, Load, Memory, Network

Other important features of PMM include:

- High Availability support via a ProxySQL Overview dashboard
- Amazon RDS and Amazon Aurora for MySQL and PostgreSQL—PMM consumes and displays CloudWatch metrics in order to provide a full picture of database activity in RDS
- External Exporters—Graph any service using PMM!

Download your free copy of PMM at percona.com/pmm.

Percona
percona.com

PRODUCT SPOTLIGHT

TREND-SETTING PRODUCTS for 2019

Pythian

Ron Kennedy, Director of Product Management

Kick AaaS Overview: How Pythian's cloud-native analytics platform solves the data silo problem

Data-driven companies understand that the traditional data warehouse isn't up to today's complex big data challenges.

Pythian's Kick AaaS was developed to answer a specific problem experienced by our clients: that of finally opening up the possibilities locked within their disparate array of data sets siloed in different databases, and governed by disconnected departments.

Kick AaaS is the cloud-native analytics platform that solves the data silo problem. It takes the hard work out of capturing, curating and preparing data for consumption by a range of users and systems. It's enterprise-wide, cloud-native data integration and sharing made simple.

The platform scales as your needs evolve. It starts with software that provides lightning-fast provisioning of data infrastructure, then builds on that foundation to enable things such as automated ETL and data science sandboxes. Kick AaaS features work together like building blocks so you can easily add the features you need or remove the ones you don't. And because it's built in the cloud—where much of your big data lives anyway—Kick AaaS grows with your needs, allowing you to continually add data sources, and enable self-service analytics and better insights using more data sources.

The keys to the success of Kick AaaS lie in its ability to work on any of the top cloud platforms (Google Cloud Platform, Microsoft Azure or AWS), its independence from any particular BI tool, its agility to add new data at business speed, and its unique ability to clean, deduplicate and unify data at scale for better insights across the organization. This means any user can access more data using any BI or visualization tool.

Find out how Kick AaaS can help you break down data silos to drive better business value.

Pythian
pythian.com

Redis Labs

Ofer Bengal. Co-Founder and CEO

With every enterprise managing and operating in the cloud, delivering an instant experience has become crucial, impacting enterprises operating in every industry including financial, entertainment, transportation, ecommerce, and more. Redis Labs is the home of Redis, the world's most popular in-memory database, and commercial provider of Redis Enterprise. Combining the vibrancy of the open source community and world-class engineering, Redis Enterprise delivers an extremely fast multi-model database that scales easily and delivers ideal functionality with great simplicity.

Because Redis is the fastest database in the industry, companies are able to deliver on consumers' expectation for instantaneous responses from applications that touch every part of people's lives. Redis combines the best of in-memory, schema-less design with optimized data structures and versatile modules that adapt to each enterprise's data needs.

Redis Enterprise is available as a fully managed database-as-a-service, in the cloud or virtual private cloud, and is downloadable for any cloud platform or private data center deployment. Powering global companies of every size and industry, Redis Enterprise is used for a wide range of use cases including high speed transactions, job and queue management, user session management, real time data ingest, notifications, content caching, and time-series data.

Redis Enterprise is a true multi-model database which offers the quickest time to market, highest levels of performance and cost efficiency to be the primary database platform for the cloud-native world where access to information has to be everywhere, all the time, and in an instant.

Redis Labs
https://redislabs.com/

PRODUCT SPOTLIGHT

RedPoint Global

Dale Renner, CEO

Consumers are more empowered than ever before and are willing to leave brands that fail to meet expectations. Providing a consistent, highly personalized experience across every customer touchpoint is what consumers expect, and has become a strategic imperative for brands, financial institutions, consumer goods, hospitality organizations, and retailers alike.

The RedPoint Global Customer Engagement Hub™ (CEH) enables business to overcome data silos, while leveraging business rules and processes to personalize information for the customer in real time. It does this by recognizing the customer in context, utilizing AI and machine learning to build a detailed customer profile, determine the next-best action, and manage delivery and intelligent orchestration of offer management. Through an open garden approach and utilization of content from NoSQL and document databases, RedPoint provides an agile, open, connected architecture that can leverage existing marketing technology and easily incorporate new technology advancements.

RedPoint CEH integrates with the vast ecosystem of marketing technologies and offers marketers continuously updated insights into how each customer transacts with your organization, including preferences, behaviors, offers, purchases, and latest interactions across every touchpoint. This enables organizations to drive higher revenue and lifetime customer value while lowering interaction costs, truly engaging the customer where, when and how each individual prefers.

RedPoint Global
www.redpointglobal.com

Revelation Software

Mike Ruane, President and CEO

OpenInsight 10 might be the most visually appealing, intuitive, MultiValue development tool on the market today. Whether you and your staff have been developing since Dick Pick was still cutting code, or you have new hires just coming out of school, OpenInsight will let you leverage your skills and expertise to create brand new applications, or update and modernize your current offerings.

Other vendors and products offer tools promising similar outcomes. With OpenInsight, you don't need to learn a new language like Java or Python. OpenInsight uses BASIC—the same programming language you're familiar with, with new functions and calls. So, the commands you know, the functions, the case-insensitivity of the programs, IF/THEN/ELSE structures instead of spaces: they're all available in OpenInsight.

Do you want to program in a green screen editor? You can. Or would you prefer a newer one, with keyword coloring, code tips, and interactive, context sensitive help? You can do that as well.

These editors, as well as Form Designers, Reporting Tools, Database Tools, are all contained in OpenInsight's IDE: Integrated Design Environment. You'll recognize the look and feel of this tool and will pick it up easily and discover it's the exact thing needed to bring your MultiValue application up to customers' current expectations. A graphical, intuitive, easy to use application.

Why learn a new language and environment, such as Python, when there's an application development tool that will give you the same end result, but leverages your existing experience and skillset? Give OpenInsight a try. You're going to love it.

Revelation Software
www.revelation.com

PRODUCT SPOTLIGHT

TREND-SETTING PRODUCTS for 2019

Robin Systems

Premal Buch, CEO

Robin is the hyper-converged Kubernetes platform for big data, databases, and AI/ML. Robin helps enterprises achieve faster realization of critical IT and business initiatives like containerization, cloud migration, multi-cloud strategy, data analytics, and cost consolidation.

Enterprise customers today have to develop custom workflows to deploy and manage each application in their big data/AI/ML pipelines and operational databases, and repeat that for each on-premise and cloud installation, leading to high cost, complexity, and time-to-value. With Hyper-converged Kubernetes, Robin is the only solution that embeds application lifecycle management into an integrated storage, network, and cloud infrastructure stack.

As a result, only Robin makes these applications agnostic of infrastructure choices, elevates every data application to a simple managed service-like experience, and automates deployment and lifecycle management of applications and data, leading to higher agility and lower cost.

For DevOps and IT architects, Robin brings agility to react to LoB and CXO asks, higher DevOps productivity, as well as infrastructure utilization, and extends Kubernetes benefits to big data and databases without new staffing or complex projects. Robin provides 1-click deploy, snapshot, clone, scale, and upgrade for big data and databases. Robin enables 1-click cloud migration for applications including data, and guarantees QoS and SLA while consolidating big data and database workloads.

Robin enables CIOs to execute a faster roll-out of critical IT initiatives (containerization, cloud migration, multi-cloud strategy, cost-consolidation) and business initiatives (AI/ML, analytics projects), while empowering the staff with self-service infrastructure and accountability.

Robin Systems
www.robin.io

Rocket Software

Julianna Cammarano, Director, MultiValue and Business Intelligence Product Marketing

Rocket UniVerse, a core component of the Rocket MultiValue Application Platform, is the backbone of critical business applications around the world. As such, UniVerse must not only meet the demands of today's applications, but also exceed expectations for performance and reliability.

The latest release of UniVerse rearchitects the core database processing engine to support a high-efficiency, high-availability, recoverable file system (RFS). UniVerse now maintains data integrity while providing the accelerated transaction throughput speeds that high-volume business applications require.

RFS, a complement to HA/DR, maintains a persistent change log (journal) so that when unexpected outages occur due to fire, flood, other natural disasters or plain-old network failure, the file system can quickly restore to the last complete transaction. RFS users will:

- Never lose a transaction,
- Reduce the risk of file corruption,
- Expedite restoration when required.

R&D test results show that high-volume ATM transaction speeds were 10% faster, and data entry processing was 60% faster. Key performance benefits include:

- Intelligent queue management combines all record updates into one,
- Field-level updates eliminates the need to write the entire record in replication, audit logging, and within your own application,
- Intelligent query optimization evaluates SQL statements and refines the order for optimal processing,
- Performance monitoring for deep diagnostics pinpoints trouble areas so you can fine-tune your system.

Rocket UniVerse not only meets the demands of modern applications, but exceeds expectations for performance and reliability for critical business applications.

Rocket Software
www.rocketsoftware.com

PRODUCT SPOTLIGHT

TigerGraph

Yu Xu, CEO

TigerGraph is the world's fastest graph analytics platform. Founded in 2012, TigerGraph is designed to unleash the power of interconnected data for deeper insights and better outcomes. TigerGraph can load 100-200 GB of data per machine per hour, while the other leading graph analytics solutions require 24 hours or more.

TigerGraph is delivering the next stage in the evolution of the graph database: the first system capable of real-time analytics on web-scale data. TigerGraph's Native Parallel Graph™ (NPG) design supports Hybrid Transactional/Analytical Processing (HTAP), offering real-time ACID transactions (OLTP) and superfast multi-hop analytics (OLAP), no matter how large or complex the dataset. TigerGraph's SQL-like graph query language GSQL provides for ad-hoc exploration and interactive analysis of Big Data. With GSQL's expressive capabilities and NPG speed, users can perform Deep Link Analytics, allowing them to uncover connections that were previously too impractical to reach or too cumbersome to express.

TigerGraph supports applications such as Artificial Intelligence (AI), machine learning, Anti-Fraud, Customer 360 and Internet of Things (IoT) to make sense of ever-changing big data.

TigerGraph's proven technology is used by customers including Uber, VISA, Intuit, Zillow, State Grid Corporation of China and Alipay. It scales up to 100,000+ queries per second across datasets with 100+ billion vertices (orders, payments, customers etc.) and over a trillion edges or relationships. TigerGraph provides pre-built data schemas, queries and algorithms you can tailor to your business needs.

Download TigerGraph Developer Edition at TigerGraph .com/developer to experience the world's fastest graph analytics platform.

TigerGraph
www.tigergraph.com

STEP OUT of everyday execution mode this May and join us for 3 days of practical advice, inspiring thought leadership, and in-depth training. Join your peers to learn, share, and celebrate the trends and technologies shaping the future of data. See where the world of Big Data and data science is going and how to get there first.

BROUGHT TO YOU BY

database TRENDS AND APPLICATIONS

BDQ BIG DATA QUARTERLY

ORGANIZED AND PRODUCED BY Information Today, Inc.

CONNECT: #DataSummit

APPLICATIONS ▸▸

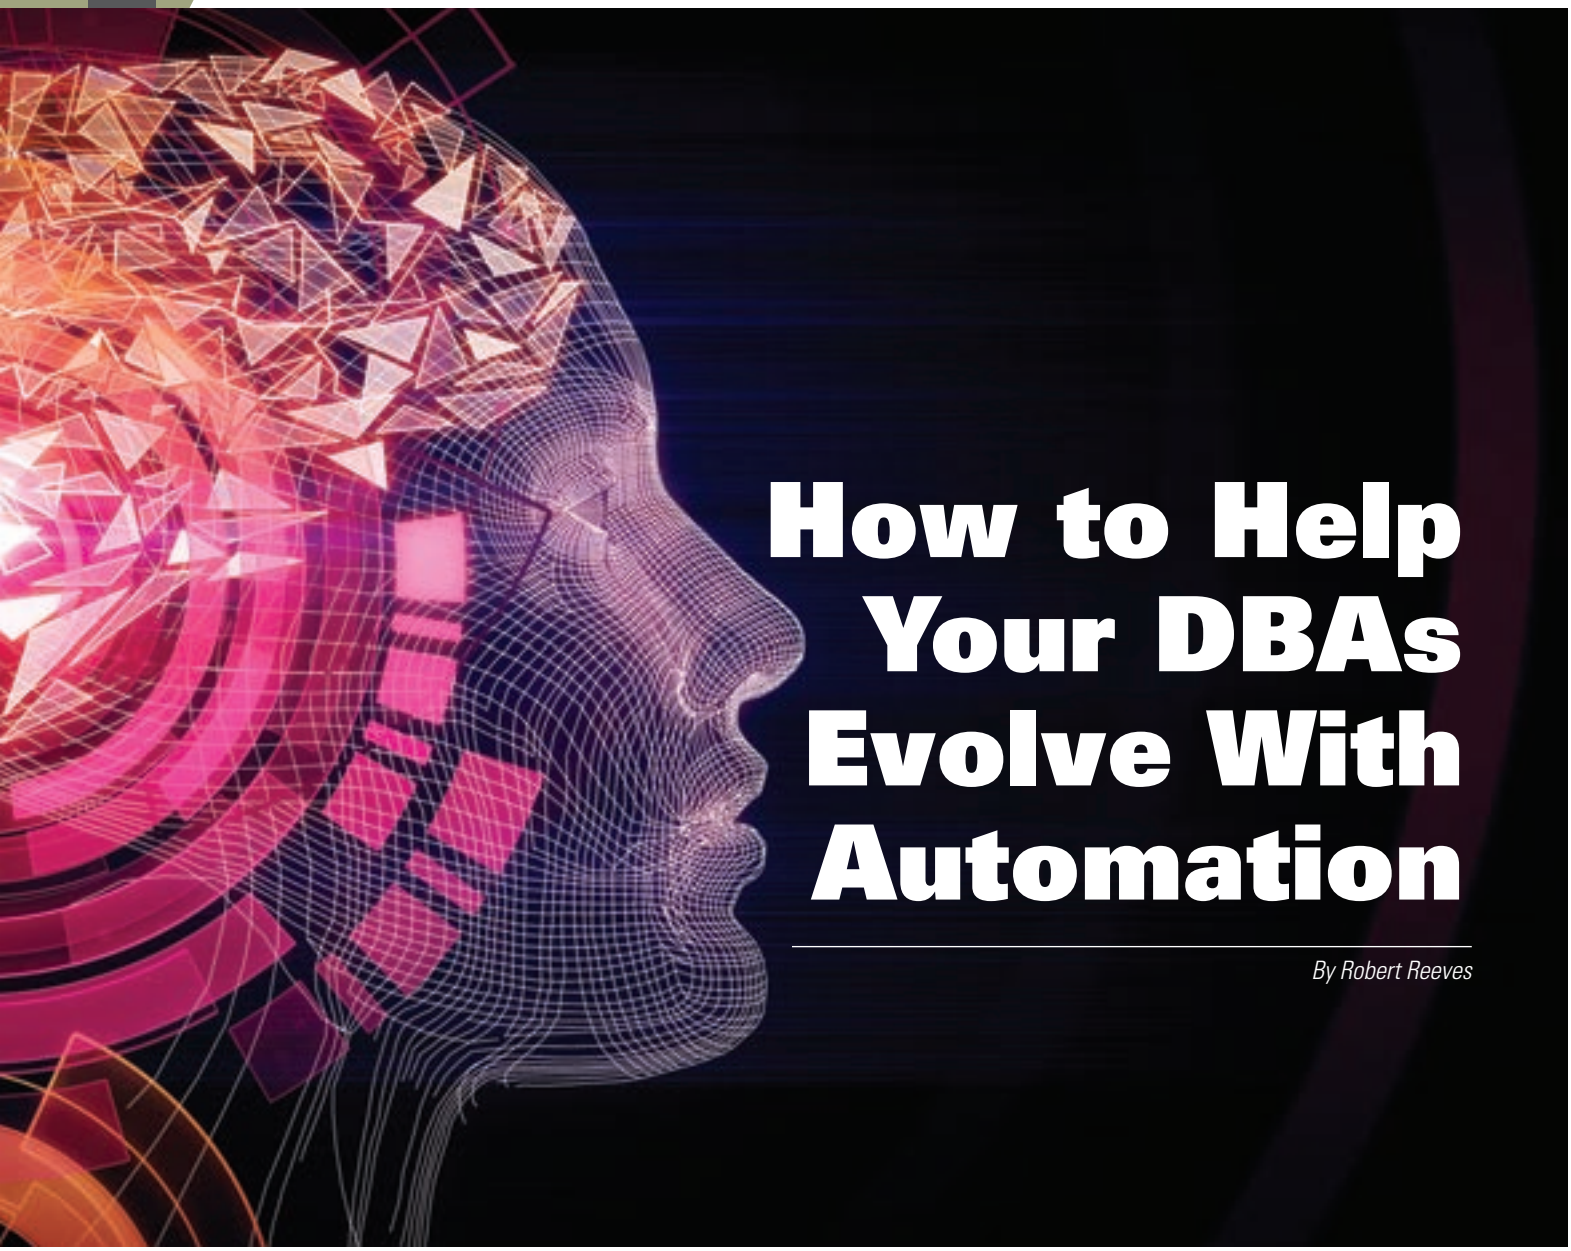

How to Help Your DBAs Evolve With Automation

By Robert Reeves

Everything changes—especially when we seek to automate tasks. Automation is leading to amazing consumer benefits—from vehicles to clothing to voice-activated devices—and making our lives better. In the workplace, as automation becomes applied to repetitive tasks that are being handled manually, people become concerned about their livelihoods. Will they lose their jobs? How will they provide for themselves and their family? How will automation impact them personally? These are natural and healthy questions. But they do not change reality, especially in the world of software development, where sooner or later every repetitive task will be a candidate for automation—even those once considered sacred and performed by the DBAs, e.g., SQL script reviews and deployments. Automation will inevitably impact everything from code development and application testing to database deployment. Evolution is needed to respond to change—and here is why that is a good thing.

Let me be clear from the start: Automation replaces tasks, not people. Automation will change how DBAs perform their tasks, but it will not change the value they can provide their company. We have seen this before with Infrastructure as Code (IaC) and tools such as Puppet and Chef. Prior to adopting IaC, system administrators automated much of their jobs, just not in a way that enabled self-service. That automation was limited to making the system administrator's life easier, not the end user's. Now with system provisioning automation, end users can request systems that meet corporate standards without having to go through a ticket system or procurement. This has provided benefits to the end user and the business but also the system administrators. The system administrator role has shifted to roles such as site reliability engineers or other more interesting and strategically beneficial roles. It's time for DBAs, and their companies, to share the same benefit.

Today, the DBA role is insular and focused on protecting data. (Data is, after all, the most valuable asset a company owns.) However, tomorrow requires data to be available and managed in a more decentralized fashion. DBAs must balance data protection with data availability to better serve application development and testing teams. Tomorrow is a "You build it, you run it" world that requires tasks such as database schema changes to be made just as easily as application code changes. Moreover, DBAs are often siloed by the database platform they support. With the advent of more specialized data platforms, being an Oracle DBA alone is not enough for some companies; they are demanding more from their DBAs.

And there's no time to waste, as the adoption rate for new technology has increased dramatically during the past century. Electricity was made commercially available in 1873, and it took 46 years for one-quarter of Americans to use it; the internet first became available in 1991, and it took just 7 years to reach the same level of adoption. The rush to containers by cloud providers and development teams is happening even faster, which makes it all the more pressing for DBAs to evolve at a pace that can keep up with these changes.

DBAs must start to think about their jobs differently. Instead of simply focusing on data control and safety (which is important), DBAs must also think about other business imperatives, such as application delivery speed, platform flexibility, and data accessibility. To balance what may seem to be competing interests, DBAs should consider evolving their practices and incorporating automation. This will help meet the demands of the business.

Automation is key to this necessary evolution. DBAs must start thinking in terms of "as a service" instead of bespoke, one-off, discrete tasks. If your only entry point to the DBA is a ticket system, then you and your company are not evolving. Self-service must be the priority. Of course, the need for control and safety has not disappeared, but the way it is applied must change. Just as IaC usage provides a menu of system options for end users, DBAs must provide options for their end users that allow them to reach their goals in the manner they deem best. This is very much similar to offering choices to a child. As my son matured and sought more screen time for video games, I didn't hand him a game controller. We started gradually offering choices such as at bath time: "Do you want to stay in the tub or get out now and play video games for a few extra minutes?" Sometimes he chose now and other times he chose to stay. Not to liken software developers to children, but they do not have the experience and resulting pattern recognition skills that a seasoned DBA has. But, the solution cannot be, "Let me see your SQL script," followed by a lengthy manual review. DBAs must balance safety with speed.

To that end, DBAs must enhance their skill set and evolve to meet these demands. Learning a programming language other than SQL should be the first step. I recommend Python, as IEEE has once again put it at the top of its list as it continues to expand its lead on C++, C, and Java. All modern tools have a Python API, and database automation tools are no exception. (If not, you're using the wrong tool.) Furthermore, experience with a database platform other than the one you are certified on is a must. Find out which platform your CTO has identified as the company's future and learn that. Today, we no longer need to sit in a classroom for 2 weeks and take a test to become certified. There are endless and, often, free resources online to take advantage of, mostly from the vendors themselves.

But this evolution will not happen in a vacuum. Today's DBAs are the most overworked group in IT. To demand that DBAs add yet another list of tasks to their already-overflowing email inbox is a fool's errand. Management must make this a priority and help. IT management must understand recent computer science grads are not becoming DBAs. Programs must be created and funded to retain DBAs and help them evolve. If not, those DBAs will find more progressive companies and take the knowledge and experience between their ears with them.

*DBAs must start thinking in terms of **'as a service'** instead of bespoke, one-off, discrete tasks. If your **only entry point** to the DBA is a ticket system, then you and your company are not evolving.*

I have seen firsthand what companies are able to accomplish when they make this change a priority. These companies can deliver compelling services to their customers far faster than their competition … and their stock performance reflects the results. Furthermore, employees are happier, more engaged, and able to deliver better results for the management team. Remember: Companies that are high IT performers deliver better stock returns than their lower-performing peers. And, that is exactly what management's compensation is tied to. Bottom line: Help your DBAs evolve and you are going to make your bonus. ■

Robert Reeves is co-founder and CTO of Datical, a provider of agile database automation solutions.

APPLICATIONS ▸▸

Blockchain Fundamentals

Q&A With Paul Tatro, Founder of Blockchain U Online

Blockchain, the distributed ledger technology, has the potential to impact a diverse range of industries from agriculture to accounting to healthcare. Supporting integrity and trust among entities, blockchain can help track the movement of goods and services—and data itself. Recently Paul Tatro, founder of Blockchain U Online, talked about where blockchain is used now and what's ahead in 2019.

What is the big advantage of blockchain?

By delivering records that convey offer and acceptance, a blockchain is a value exchange protocol that provides a trust layer for the internet and then digitally records the data in a shared distributed ledger in packages called blocks. If you look at the idea of "offer and acceptance," that is a fundamental element of any transaction so that gives you the idea of the ubiquitous potential of blockchain because it can be involved anywhere that you have an offer-and-acceptance type of scenario.

What is the significance of a "value exchange protocol"?

Through a blockchain, you are able to not only say, "We are going to do a deal together" but also settle on the terms, conditions, and the value of that deal as part of the transaction being completed by the blockchain. The concept of a value exchange protocol is very powerful and this is where this idea of eliminating trusted third parties in the middle of a transaction—that is where the possibility comes from.

And the "trust layer for the internet"?

The internet is becoming almost like electricity. If you don't have it, you can't function; but we also know that the internet is full of traps, so to have something that can put trust into an application that is available in a distributed fashion across the internet is something that is very unique and powerful, and it is what gives blockchain its appeal.

What are the key benefits of blockchain for the enterprise?

First off, it provides the highest degree of accountability in any application environment. This is where aspects such as consensus and the value exchange protocol come in. These are part of the accountability of the application and so third parties in the middle are not needed.

Blockchain also guarantees the validity of transactions by recording them in several places. It is a distributed type of application so there isn't a single point of failure. If there are 100,000 nodes participating in a blockchain, then there are 100,000 copies of the transaction, so to speak. It makes it more difficult for hackers, and the integrity of what happened is guaranteed because it is recorded in so many different places.

How are the other advantages?

Because you are not relying on someone looking over forms and deciding what is valid, you have effectively taken the process for validation and codified it so the process integrity is ensured. The computer software will run the same every time. The other thing is the immutability of it. It is virtually impossible to make changes to a blockchain once a block has been added. That ensures that if data is in a blockchain, it has not been changed. When you combine the immutability of blockchain with the resistance to hacker attacks due to its distributed nature, I think it yields the highest degree of integrity that you will find with any application.

What are some of the areas that need to be strengthened in order for blockchain to take a more prominent place in the enterprise?

There are a couple of things that have to get sorted out before people jump in with both feet. One is the idea of regulations. Until governments decide how a truly distributed application is enforced or governed, it is going to be tricky. Let's say there is a blockchain spanning jurisdictions and geographic boundaries—such as Bitcoin. Which rules apply? Should it be U.S. rules? Maybe there was an issue that took place in Dubai. Who really is involved? There is an old saying, "When more than one person is in charge, one is to blame." That needs to be clarified because companies don't want to find themselves vulnerable to laws they didn't anticipate. This could happen with a public or private blockchain due to the distributed nature.

What else?

A second concern is standards. There need to be standards about how blockchains work. Currently, there are more than 25 different consensus methods in blockchain technologies. Each one works differently because of the consensus method. Bitcoin uses Proof of Work. Ethereum used to use Proof of Work but is shifting to a specialized Proof of Stake called Casper. EOS has a specialized Proof of Stake called Delegated Proof-of-Stake. There are other protocols as well, such as Proof of Burn, and it goes on and on. To me, this adds to the confusion around blockchain. The concepts are straightforward, but when you get into some of the details, it gets to be a confusing ball of yarn. This is where some standardization needs to come into play.

How widespread is the use of blockchain in the enterprise now?

In corporate America, everyone is kind of looking at it, but I read a McKinsey article that estimated that 90% of the blockchain projects being undertaken today will never see the light of day. In other words, companies are learning, experimenting, and trying to figure out what the impact will be and what the process will be.

Obviously, one of the key things coming out of blockchain is an alternative currency like Bitcoin. Having an alternative currency has a greater potential impact than many people realize, especially when you understand that more than 70% of the world is unbanked. Bitcoin, particularly, has the potential to bring those people in as participants to the worldwide economy, and that is the big buzz around those kinds of currencies.

Are there other areas?

There was an announcement recently by IBM and Maersk about a blockchain shipping solution, so that is an example of another area. I think identity management is going to be yet another area. There are companies, such as one called uPort, that are advancing the concept of self-sovereign identity so instead of some big company controlling who looks at an individual's information and controlling what information is dished out, the individual can control who looks at their information and what information is provided. The notion gets more interesting when you consider attestations of who you are and what you have done.

What do you see on the horizon for 2019?

We are effectively in what I call version 3 of blockchain. Version 1 was basically Bitcoin, a simple blockchain. The concept of smart contracts was not a part of it, and it did one thing—it transferred Bitcoin between parties. Version 2 was Ethereum, which introduced the first smart contract language that enabled business logic or rules to be put into a program and stored in a blockchain.

What is version 3 of blockchain?

This is where additional services over and above smart contracts are becoming part of the blockchain. This will allow corporations to see what security and database components are available. These are things that people have taken for granted in the evolution of information technology over the years, and are now coming to blockchain. Governance over what happens in a blockchain environment is another huge issue that needs to be considered. An additional aspect is the network factor. If a company wants to introduce a shipping application, in a blockchain world, it has to get all the players involved. The dock workers, the loaders, truck drivers—all the companies involved—have to become participants in the blockchain network for it to work effectively. So, to me, this idea of the network effect will be one of the biggest challenges. ■

This interview was conducted, edited, and condensed by Joyce Wells

MV SOLUTIONS ▸▸

MULTIVALUE AND THE CLOUD: FLEXIBILITY FOR THE FUTURE

By Julianna Cammarano

With the ever-changing business landscape, many companies that rely on MultiValue platforms are looking for ways to integrate existing systems with new technologies—and specifically, figure out how to leverage the capabilities of the cloud.

Organizations often find themselves torn between the platforms they've built their critical business applications on and user demands for modern, intuitive interfaces that provide anytime, anywhere access to web and mobile applications.

In today's evolving technology space, we have arrived at a tipping point. It is easier than ever to embrace the range of capabilities by hosting application development platforms such as MultiValue in the cloud. This approach can deliver value to organizations in many different ways, including improved performance, more overall business elasticity, and new levels of automation as well as support for DevOps.

To be fair, replacing a MultiValue system can be expensive, risky, and time-consuming, which is why companies that depend on MultiValue platforms sometimes feel limited by their options. But now is a perfect time for companies who rely on the power of MultiValue to modernize and rejuvenate existing applications by moving them to the cloud.

For starters, the cloud today provides a low barrier to entry for organizations looking to access powerful capabilities associated with MultiValue. If a company moves its application, it can now quickly and easily use tools that deliver a wide range of functionality and are resident in the cloud. Examples include AWS CloudWatch or Azure Application Insights for monitoring and setting up alerting rules around critical business functions. High availability and disaster recovery, required for business continuity, is also much easier to set up in the cloud.

Many organizations are achieving dramatic improvements in overall performance by deploying MultiValue in the cloud. This approach makes available a wide range of services such as Azure's Application Insights for Logging. Application Insights lets users open a dashboard and quickly identify what's running in the cloud, making it much easier to monitor activity as well as search for and extract useful information and insights.

*The cloud provides a **low barrier** to entry for organizations looking to access **powerful capabilities** associated with MultiValue.*

Another key advantage of hosting MultiValue in the cloud is the level of overall business elasticity that it provides. Organizations now have the freedom to "pay as you go," reducing the need to purchase and support a complex system that requires a major capital expense and could represent increased risk. MultiValue in the cloud frees up funding that can then be spent on more strategic business imperatives.

We are seeing an increase in running DevOps in the cloud. Businesses can now use cloud-based orchestration tools to bring up test machines, deploy a new version of an application to a small group, and then switch over to production once it's ready for general availability.

MultiValue in the cloud also provides the ability to employ automation for areas such as load balancing. For example, a retail customer has been using cloud-friendly products to stay connected to his hybrid architecture, which is a mix of on-premise and cloud-deployed instances.

Plus, your MV application is always available, even during a maintenance event. Recently, a large financial services firm tested accessibility of its cloud-based MV application during just such a maintenance session. Despite running repeated replication processes, pausing, and then re-directing traffic to the replicated server, the MV instance worked flawlessly without any issues.

Simply put, MultiValue in the cloud provides peace of mind. With an MV application development platform in the cloud, companies don't have to worry about running out of disk space or data being deleted or whether they are running the most up-to-date version of an application. It is all managed behind the scenes.

The ultimate goal, of course, is to make it easier for businesses to modernize and grow by leveraging current technologies while also building in flexibility for the future. ■

Julianna Cammarano is director of product marketing for business intelligence, analytics, and the MultiValue Application Platform at Rocket Software (www.rocketsoftware.com).

▸▸

United States Postal Service | Statement of Ownership, Management and Circulation | (Requester Publications Only)

1. Publication Title: Database Trends and Applications. 2. Publication Number. 16-230. 3. Filing Date: 10/1/18. 4. Issue Frequency: Bi-Monthly; Dec/Jan., Feb/Mar., Apr/May, June/Jul., Aug/Sept., Oct/Nov. 5. Number of issues Published Annually: 6. 6. Annual Subscription Price: 0. 7. Complete Mailing Address of Known Office of Publication: Information Today, Inc., 143 Old Marlton Pike, Medford, Burlington County, NJ 08055. 8. Complete Mailing Address of Headquarters or General Business Office of Publisher: Unisphere Media a Division of Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055. 9. Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor: Publisher: Thomas Hogan, Jr. Group Publisher, 143 Old Marlton Pike, Medford, NJ 08055-8750. Editor: None. Managing Editor: Joyce Wells, 121 Chanlon Road, New Providence, NJ 07974. 10. Owner: Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055, Thomas H. Hogan, 143 Old Marlton Pike, Medford, NJ 08055, Roger R. Bilboul, 22 Earls Terrace, London W8 6LP, England. 11. Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 percent or More of Total Amount of Bonds, Mortgages or Other Securities If none, check box: None. 12. Has Not Changed. 13. Publication Title: Database Trends and Applications. 14. Issue Date for Circulation Data Below: Oct/Nov. 2018. 15. Extent and Nature of Circulation: a. Total Number of Copies (Net press run): Average No. Copies Each Issue During Preceding 12 Months, 4,846; No.Copies of Single Issue Published Nearest to Filing Date; 4,370. b. Legitimate Paid and/or Requested Distribution. (1) Outside-County Paid/Requested Mail Subscriptions Stated on Form 3541 (Include direct written request from recipient, telemarketing and Internet requests from recipient, paid subscriptions including nominal rate subscriptions, employer requests, advertiser's proof copies, and exchange copies): Average No. Copies Each Issue During Preceding 12 Months, 4,654; No.Copies of Single Issue Published Nearest to Filing Date, 4,242. (2) In-County Paid/Requested Mail SubscriptionsStated on Form 3541 (Include direct written request from recipient, telemarketing and Internet requests from recipient, paid subscriptions including nominal rate subscriptions, employer requests, advertiser's proof copies, and exchange copies): Average No. Copies Each Issue During Preceding 12 Months, 0; No.Copies of Single Issue Published Nearest to Filing Date, 0. (3) Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Paid or Requested Distribution Outside USPS: Average No. Copies Each Issue During Preceding 12 Months, 0; No.Copies of Single Issue Published Nearest to Filing Date, 0. (4) Requested Copies Distributed by Other Mail Classes Through the USPS (e.g., First Class Mail): Average No. Copies Each Issue During Preceding 12 Months, 0 ; No. Copies of Single Issue Published Nearest to Filing Date, 0. c. Total Paid and/or Requested Circulation (Sum of 15b (1), (2), (3), and (4): Average No. Copies Each Issue During Preceding 12 Months, 4,654; No.Copies of Single Issue Published Nearest to Filing Date, 4,242. d. Non-requested Distribution (By Mail and Outside the Mail) (1) Outside County Non-requested Copies Stated on PS Form 3541. (Include Sample copies, Requests Over 3 years old, Requests induced by a Premium, Bulk Sales and Requests including Association Requests, Names obtained from Business Directories, Lists, and other sources): Average No. Copies Each Issue During Preceding 12 Months, 0; No.Copies of Single Issue Published Nearest to Filing Date, 0. (2) In-County Non-requested Copies Stated on PS Form 3541 (include Sample copies, Requests Over 3 years old, Requests induced by a Premium, Bulk Sales and Requests including Association Requests, Names obtained from Business Directories, Lists, and other sources): Average No. Copies Each Issue During Preceding 12 Months, 0; No. Copies of Single Issue Published Nearest to Filing Date, 0; (3) Non-requested Copies Distributed Through the USPS by Other Classes of Mail (e.g., First-Class Mail, Non-requestor Copies mailed in excess of 10% Limit mailed at Standard Mail or Package Services Rates): Average No. CopiesEach issue During Preceding 12 Months, 0; No. Copies of Single issue Published Nearest Filing Date, 0. (4) Non-requested Copies Distributed Outside the Mail (Include Pickup Stands, Trade Shows, Showrooms and Other Sources): Average No. Copies Each Issue During Preceding 12 Months, 96; No.Copies of Single Issue Published Nearest to Filing Date, 25. e. Total Non-requested Distribution (Sum of 15d (1),(2), and (3)): Average No. Copies Each Issue During Preceding 12 Months, 96; No.Copies of Single Issue Published Nearest to Filing Date; 25. f. Total Distribution (Sum of 15c and e): Average No. Copies Each Issue During Preceding 12 Months, 4,750; No.Copies of Single Issue Published Nearest to File Date; 4,267. g. Copies not Distributed: Average No. Copies Each issue During Preceding 12 Months, 96; No. Copies of Single Issue Published Nearest to Filing Date, 103. h. Total (Sum of 15f and g): Average No. Copies Each Issue During Preceding 12 Months, 4,846; No.Copies of Single Issue Published Nearest to Filing Date, 4,370. i. Percent Paid and/or Requested Circulation (15c divided by f times 100): Average No. Copies Each Issue During Preceding 12 Months, 97.97%; No.Copies of Single Issue Published Nearest to Filing Date, 99.41%. 16. Publication of Statement of Ownership for a Requestor Publication is required and will be printed in the December 2018/January 2019 issue of this publication. 17. Signature and Title of Editor, Publisher, Business Manager or Owner: John C. Yersak, Vice President & CAO. Date: 10/1/18. I certify that all information furnished on this form is true and complete. I understand that anyone who furnishes false or misleading information on this form or who omits material or information requested on the form may be subject to criminal sanctions (including fines and imprisonment) and/or civil sanctions (including civil penalties).

MV SOLUTIONS

jBASE Helps Encompass Supply Chain Streamline Development Operations

Zumasys, a provider of NoSQL database software for business-critical PICK applications, has announced that Encompass Supply Chain Solutions, Inc. has migrated its custom ERP system to jBASE.

jBASE, Zumasys' flagship product, is a NoSQL database developed nearly 30 years ago and used by the largest international banks in the world.

Encompass also transferred its infrastructure from a co-location facility to the Zumasys cloud, now owned and operated by NexusTek.

The move has helped Encompass eliminate costly downtime and streamline development for the company's primary business applications.

With its ERP application running on jBASE, Encompass can now take advantage of improved reliability and performance for batch processing and web apps.

As Encompass continues to grow its business and service offerings, it is critical to have the right systems in place to run its operations and scale with confidence, said Encompass president and CEO Robert Coolidge. "We are very pleased with Zumasys' ability to develop a strong, dependable platform to provide exceptional uptime and support for our customers."

According to Encompass, its business success depends on its custom-developed ERP application, which feeds off its PICK MultiValue database platform. As the company's database approached its end of life, it was no longer being actively supported or updated. Over time, the database became slow and unreliable, resulting in downtime for the company's ERP system and ecommerce applications.

Encompass ruled out moving to SAP or Oracle, as switching database platforms would have been too costly, risky, and time-consuming. The company also considered additional MultiValue database solutions before deciding on jBASE from Zumasys.

With jBASE, Encompass could retain its custom features with a clear path for the future. Zumasys also offered a truly cloud-ready database solution. "By moving to the cloud, we can now focus our resources on development instead of running hardware," added Brent Blair, Encompass VP of IT. ■

ROB MANDEVILLE

Rob Mandeville is a senior DBA at SolarWinds, experienced with major RDBMS vendor engines.

Anomalies—Predicting the Past

Definition of anomaly: something that deviates from what is standard, normal, or expected.[1]

If you are a modern database professional, you have likely heard about or looked into anomalies and how to detect them. At a very simple level, anomaly detection looks at historical values and predicts future values based on them. An anomaly occurs when we miss the prediction with enough significance.

In a very literal sense, we are predicting the past.

Innate Anomaly Detection

We come with an innate ability to pick out anomalies. In fact, part of our success as a species can be attributed to anomaly detection. The idea in this case is to identify outliers for more rigorous inspection or analysis.

- That spider doesn't look similar to other ones I've encountered.
- The branch wasn't broken before.
- This creek is much lower than usual for this time of year.

Noticing these things can help with survival. It could be argued that's a good thing. So, the fact that one of the hot focuses of data science is anomaly detection should come as no surprise. However, anomalies by themselves are not good or bad; they are just different from the norm. No judgment. Additionally, the absence of an anomaly doesn't imply good or bad—it just implies normal. There still has to be the human element to help define the good and bad.

The Science

If we are going to try to automate anomaly detection programmatically, we need to decide on an algorithm that fits best. This is likely an algorithm that best predicts future datapoints based on historically observed datapoints and gets better over time with more data. The solution would also likely account for seasonality (same day of the week, same month of the year, etc.). On the surface, anomaly detection sounds great; and it can be, as long as you don't base too much judgment on it without the investigation and analysis as to why the anomaly occurred.

Before we get to judgment, there's still work to be done—we're not quite done with how we detect anomalies with our algorithm yet. We have to define how far off from expected the new data point is to call it significant (standard deviation being a popular choice). Once we have put some math behind defining how our algorithm appears and how far off any next value needs to be before we call it abnormal, we can start looking for anomalies in an automated fashion.

Example

Consider two workers—both of their jobs have them starting at 8:00 a.m. One always arrives a few minutes before 8:00 a.m. and one always arrives right at 8:30 a.m. As long as their arrival patterns do not vary, there will be no anomaly. If the 8:00 a.m. worker arrived at 8:05 a.m., would we call that an anomaly? Or how about 8:15 a.m. or 8:30 a.m.? What if the second worker arrived at 8:00 a.m. one day? 7:30 a.m.? Would we call those anomalies even though the behavior is better than the norm? In the strictest sense of anomaly detection, we would call everything out that was considered an outlier. Then, depending on whether the outlier was below or above our expected value, we would be ready to place a judgment.

The Judgment

At this point, we are ready to assign a good or bad determination to our anomaly. In our example, the second worker comes in late each and every day. The fact that one day she came in at 8:00 a.m. or earlier was a good thing. In fact, it was the normal that was bad. Now measuring something this simplistic can of course be corrected by also comparing data values with an absolute or threshold, but just using anomaly detection, this might create a flawed model. In fact, check out this post (www.dbta.com/Columns/Next-Gen-Data-Management/NEXT-GEN-DATA-MANAGEMENT---Dangers-of-Statistical-Modeling-127922.aspx), where I discuss statistical bias and how it can impact results in quite an unintentional, yet negative, way.

The Wrap Up

Using anomaly detection to surface deviations from the norm can be a powerful tool when automating monitoring for just about any statistical value. However, it will not uncover the bad behavior that may be happening all the time and is only part of the story. Combining data from anomaly detection *plus* data that can surface good behavior versus bad behavior when it occurs as the norm starts to get into more intelligent analysis. The hype is legit. Anomaly detection can be a very powerful tool—just not in complete isolation.

I did not cover the concept of confidence interval intentionally—saving that for a later date. ■

[1]*Cited from Google Dictionary using search term anomaly definition*

GUY HARRISON

Guy Harrison, a software professional with more than 20 years of experience, is a partner at Toba Capital and the author of *Next Generation Databases* (Apress). Contact him at guy@tobacapital.com.

Web Services Move Forward With GraphQL

It's gratifying to be in the business long enough to see new paradigms transition from lofty visions to everyday standards. Such is the case with web services. When I first started writing a column for *DBTA*, the evolution of web services was well underway, although a battle for dominance between competing standards was incomplete. In particular, there was competition between a set of committee-defined standards (WS-this and WS-that) and a looser set of techniques based around simple protocols.

EMERGING TECHNOLOGIES

As is so often the case in the software industry, developers continued to program while the standards bodies deliberated. By and large, programmers rejected standards-based mechanisms in favor of a far more flexible approach: implementing web services by transmitting XML documents across HTTP. Programmers discovered that they could use HTTP directives and URL constructs to eliminate a lot of the complexity required by SOAP and its cousins.

Eventually a widespread standard emerged for this simplistic approach: REST (Representational State Transfer).

Strictly speaking, REST is an architectural style, not a platform. To be "RESTful," an API must adhere to several core principles first articulated by Roy Fielding. These principles define how a web application navigates through network of webpages or web resources. REST defines how requests identify resources and navigate to related resources. For instance, a REST request might first identify a specific customer account, then use the results from that request to create a request that navigates to that customers profile.

REST gained popularity throughout the 2000s, and today, REST using HTTP and JSON has created an almost universal language for web APIs. In so doing, REST has done more than any other technology to bring the vision of web services to reality.

However, similar to all practical general-purpose technologies, REST is not perfectly suited to all scenarios, and there have been ongoing attempts to design a "better than REST" platform for web services.

GraphQL has emerged as the favorite alternative to REST for modern web API design. GraphQL was first used internally at Facebook before being open sourced in 2015. GraphQL is described as a data query and manipulation language which at first glance might suggest it has more in common with SQL than with REST. However, in reality, both GraphQL and REST are applicable across a very similar range of web API scenarios.

GraphQL implements a friendlier syntax than REST and provides significant efficiency gains. REST calls typically return all data related to a resource, while a GraphQL query can request subsets of data, reducing network overhead.

GraphQL also allows a request to return related data without having to create specialized API endpoints. For instance, a GraphQL request could request details of a customer and all of that customers' orders, implicitly navigating the relationships between those entities in the server. In REST, such a request would have to be anticipated and explicitly created or would have to be resolved using multiple REST calls.

GraphQL also provides a strongly typed schema which allows clients to navigate the API endpoints more effectively. This allows for ad hoc requests and intelligent API browsers.

GraphQL is unlikely to completely overturn REST but appears to take web services one further step forward toward maturity. GraphQL is likely to further expedite the transition from monolithic applications to loosely coupled microservices.

The transition to microservices is generally regarded as a good thing. However, there is a risk that tomorrow's unmaintainable applications will be a confusing mess of poorly defined microservices. We'll need strong, self-describing protocols to avoid this scenario, and GraphQL seems a step in the right direction. ■

KEVIN KLINE

Kevin Kline, a longtime Microsoft SQL Server MVP, is a founder and former president of PASS and the author of *SQL in a Nutshell*. Kline tweets at @kekline and blogs at https://blogs.sentryone.com/kevinkline.

Loads of Data and AI Announcements at Ignite 2018

The 2018 Microsoft Ignite conference was overflowing with attendees this year, as user enthusiasm continues to grow with the advent of CEO Satya Nadella. Since space is short and the announcements are many, let's get straight to the details of the product innovations available for data professions.

New in AI

If you're interested in voice and natural language processing, then here are a couple of announcements just for you:

- Not long ago, Microsoft had several distinct artificial intelligence (AI) speech capabilities, including speech recognition, speech translation, and customized models to create a unique voice for your app. Those have been combined into Speech Service, now generally available. In public preview is Human Parity Text to Speech, which utilizes the latest in deep neural network technology to make computer voices nearly indistinguishable from real humans. Details are at **https://bit.ly/2PtWT5r.**
- Furthering the goal of making your apps behave even more similar to that of a human being, the Microsoft Bot Framework v4 SDK is now generally available. You can program in C#, Java, Python, and JavaScript. It includes many tools for simplifying and building bots with a modular and extensible architecture, allowing easy selection of the components and services you need. Details are at **https://bit.ly/2OSpJwR**.

SQL SERVER DRILL DOWN

- Want your apps to include a conversational personal intelligent assistant (PIA) that knows how to complete specific corporate tasks? Check out the new Cortana Skills Kit for Enterprise. Underpinned by Azure Bot Services and Azure Language Understanding, imagine scenarios such as building a Cortana skill that lets employees ask the PIA to schedule their vacation time, saving 30 to 60 minutes of intranet surfing to figure out how to make the arrangements themselves. This program is by invitation only and not all features have been announced. See **https://bit.ly/2ArAIEz** for more information.

My personal interests in AI mostly focus on machine learning. Here, Microsoft further advances its machine learning capabilities and also puts more AI into its own products, such as in Microsoft 365, Excel, and Dynamix 365.

- New in Azure Machine Learning is the automated machine learning feature set which enables users to automate: data transformations at speed, model selection for more efficient algorithms, and hyperparameter tuning to quickly learn the accuracy of a given pipeline's predictions. The new SDK for Python includes features for distributed deep learning, allowing massive clusters of GPUs, access to field programmable gate arrays (FPGA) for amazing speed at image processing, and easier integration with IDEs such as Visual Studio Code, Jupyter notebooks, Azure Databricks notebooks, and PyCharm. Read more at **https://bit.ly/2PSp2Qp**.

New in the Data Platform

Microsoft used Ignite 2018 to announce the public preview of SQL Server 2019, with a special focus on big data features. First, Spark and Hadoop distributed file system (HDFS) are now built into SQL Server to accelerate ingestion, storage, and analysis of data at petabyte scale. There are also new direct query connectors to Oracle, Teradata, and MongoDB. As with every release, there are many security, usability, and optimization improvements. But I will cover those in a separate, more detailed article.

Details are at **https://bit.ly/2CFG8x1**.

Not long ago, Microsoft had several distinct artificial intelligence ***speech capabilities****, including speech recognition, speech translation, and customized models to create a unique voice for your app. Those have been combined into* ***Speech Service****, now generally available.*

Azure SQL Database offers improved query performance features under the *Intelligent Query Processing* moniker. This feature set includes row-mode memory grant feedback, approximate query processing, and table variable deferred compilation. All three features work together to make SQL processing faster, more responsive to memory usage issues, and improve long-standing SQL coding issues. And, as happens with every release, Microsoft has announced more pricing and performance tiers for great flexibility (and potentially lower cost) when deploying apps to the cloud, such as the new Azure SQL Database Hyperscale, a way to deploy a single but highly scalable database that can grow from a few to hundreds of terabytes in size. Azure SQL Data Warehouse also provides a new, lower-entry point to help customers get started more quickly. Details are at **https://bit.ly/2NKNOVw**.

Microsoft's NoSQL offerings have also grown:

- The Spark-like Azure Databricks product now provides Azure Databricks Delta as public preview, improving data reliability, simplified data pipelines, and improved job and query performance. More information is at **https://bit.ly/2NRBtz5**.
- Azure Cosmos DB provides multi-master support for high availability with single millisecond latency and better conflict resolution. The Reserve Capacity feature reduces costs for using Cosmos DB. And the Cassandra API is further enhanced for users familiar with that NoSQL platform. Details are at **https://bit.ly/2RgbMoR**.
- Azure Data Explorer is in public preview. This is a new indexing and querying service for interacting, lightning-fast ad hoc data exploration from data that originates in apps, servers, and edge devices. Details are at **https://bit.ly/2yuQPzN**.

Other News for Data Professionals

There's more than one way to learn what's new, starting with the new Microsoft Learn and the role-based Microsoft Certifications Microsoft Learn. This is free, interactive web-based training with a tutorial approach to teach Azure and Business Applications. There is cool progress tracking and gamification features such as achievements, while using free Azure resources for hands-on learning. There's also a new role-based set of Microsoft Certifications aligned to job roles.

Go to **https://bit.ly/2N8hGWL** to learn more. ■

SIMON PANE

Simon Pane is a principal consultant working for Pythian and an IOUG board member. As an Oracle ACE, he is a regular writer on the Oracle blog (https://blog.pythian.com/author/pane), an editor for IOUG SELECT Journal, and frequently speaks at major Oracle conferences.

The Tools the Modern DBA Needs to Know

I often work with student or junior DBAs who ask about the skills they will need to be effective in the modern technical landscape. And, as the technology changes, of course so do the required skill sets.

Sometimes I hear questions such as: "Should I learn Python or PHP to become an effective DBA?" My answer is usually, "Neither." And to expand, I usually recommend that they don't focus on languages but rather on tools. I don't mean learning the intricacies of tool commands—those can be looked up—but more importantly tool concepts and fundamentals.

The next logical question is about which tools the modern DBA needs to know. In my opinion, the tool landscape that the DBA needs to understand is moving toward the DevOps realm. That doesn't mean that DBAs need to be DevOps experts—they're still different roles. But rather, there is some overlap and, at the very least, the DBAs need to be familiar with and, in some cases, use these tools.

In the remainder of this article, I'll explore the tools I personally think are most relevant to today's DBA.

Ansible

If there were only one DevOps tool that the modern DBA should know, it would be Ansible. Often, I see DBA tasks that need to be performed against multiple databases in the environment. Maybe against all instances in the estate, maybe against just production or non-production, or maybe against a named group of related databases. When environments are large, I've seen a "divide and conquer" approach where a DBA team might split up the work, saying that person A will handle this group of database; person B, this other group; and so on. However, Ansible is a much better solution for these situations.

Whenever the modern DBA needs to execute the same tasks against multiple databases, the answer should always be Ansible. Ansible is ideal for "automating manual tasks." It is a free scripting language/tool but part of the beauty is that it's completely agentless—meaning that the Ansible software itself is only installed on one central computer such as a shared DBA machine or even the DBA's desktop/workstation. From there, it performs all of the remote commands against target machines via Secure Shell (SSH) protocol. There's no software to install on the database servers—all that's required is connectivity via SSH. In addition, targets can be easily grouped into an unlimited number of permutations. For example, "All"; "All prod"; and "All application X"—whatever is required.

There are many online resources and videos to help DBAs get started with Ansible. Once they see it in action and understand its simplicity and power, DBAs usually start to quickly see opportunities for how to use it in their day-to-day jobs.

Git

The next logical tool that DBAs should understand is the concept of Git. Typically, this is in the form of GitHub (the largest online Git service, now owned by Microsoft). But GitHub isn't the only Git implementation. GitLab, Atlassian Bitbucket, and local Git implementations are other similar options.

The idea behind learning Git for the DBA is script management. Recognizing that administrators are not developers, they will always have some number of scripts which they've developed or must maintain. And maybe/hopefully, some of those are Ansible scripts!

Traditionally, DBAs would store these in a shared file system or maybe an NFS mount or even a Windows shared drive. However, in the modern world, DBAs should be storing all of their tools and scripts in a Git repository such as GitHub. This provides a central repository, change history, and version comparisons, as well as the ability to easily share externally if desired (repositories can be private or shared—functionality and costs depends on the service used). DBAs typically use Git in a simplified workflow, although proper development teams will use it in a more sophisticated manner that includes forks, branches, and other such advanced features.

In fact, DBAs are probably already recognizing that GitHub has become the standard location for external sharing of DBA tools and utilities. It is likely that their favorite public domain tuning script is already being hosted in and shared via GitHub. The beauty of DBAs learning Git is not just that it allows them to do their own internal script management, but that it also allows them to provide feedback, suggestions, and updates to these public domain tools. ▶▶

Documentation Systems

Another important recommendation from my perspective is to use a proper wiki-style documentation system. If you're still documenting run books and other how-to types of documentation in Microsoft Word (and/or Excel), then I'm afraid you're already a dinosaur. Such documents are very difficult to keep up-to-date, version properly, and even find. The solution is a proper wiki-style documentation system.

My personal preference is Atlassian Confluence. It's a very easy-to-learn WYSIWYG wiki-ish system which allows for easy editing, collaboration, and sharing. With a little CSS customization, it even does a reasonably decent job of exporting content to PDF.

Even if Confluence isn't an option for you, there are other wiki tools, including ones that are installed locally and others that are online services. Regardless, the point is that any documentation including run books and how-to manuals should be in a system that is easy to update, provides revision history, is easily shared, is easily searchable, and ideally also provides collaboration (i.e., commenting functionality). Document editors and spreadsheets do not meet these requirements.

Vagrant and Docker

Often DBAs want "test systems." It's easy to provision virtual machines either with or without software pre-installed online from cloud providers, but sometimes administrators are afraid of cloud costs or for other reasons want to do this locally.

So how can we bring some of the ease and speed of provisioning that we see with online cloud systems to our local environment? The answer is tools such as Vagrant and Docker. Both allow us to quickly spin up virtualized environments with consistency and automation. But which to use?

Deciding which is best, in my opinion, depends on your needs. Docker is based on the concept of containers and may be best when you need multiple copies of similar systems. For example, maybe you need eight environments that are almost exactly the same so you can test scalability, replication, or failover options between them. In a case such as this, Docker is quick and efficient—and usually using containers is significantly more resource-efficient on host systems than full virtual machines are.

But in other cases, DBAs need a mix of environments. They may need to easily create (and soon destroy and re-create) a variety of systems such as different flavors of Linux alongside a Solaris system and a Windows server so they can test how their platform-agnostic tools or scripts will work against all of these. In cases such as these, I tend to recommend Vagrant. Vagrant could almost be considered "a scripting tool for automated deployments of virtual machines." This means that Vagrant allows for the easy creation and destruction of VM environments, including those for Oracle Virtualbox. Vagrant can start with publicly available VM images (called "boxes") or you can create your own custom ones as a starting point. From there, Vagrant scripts can be used to customize the machines upon creation.

Other Honorable Mentions

Terraform

If DBAs are responsible for provisioning cloud resources, then DBAs should also familiarize themselves with Terraform. Terraform is from the same maker as and has some similar concepts as Vagrant. However, Vagrant is for the provisioning of virtual machines whereas Terraform is for the provisioning of other infrastructure using code.

In my experience though, provisioning of cloud infrastructure is usually handled by system administrators, site reliability engineers, or DevOps teams and less commonly by DBAs. Consequently, it's good for DBAs to familiarize themselves with Terraform, but they should leave the actual implementations to those other experts.

Puppet

I sometimes hear comments such as, "DBAs can automate their software installations and database creations with Puppet." While technically true and I've seen this done in real life, it's not really the best tool for the job.

Puppet is best for enforcing consistent state. For example, for ensuring that a Linux iptables file is configured as per internal standards. If anyone changes the file, Puppet will revert it back. Unlike Ansible, Puppet also requires some additional components: a main server and agents on target machines.

So, for many reasons, I think Puppet is best left in the domain of system administrators. It is definitely useful but less so for DBA purposes. Most similar DBA-type activities that people may consider using Puppet for can often be implemented instead with Ansible.

Many Options

Modern DBAs should be excited to see how easily they can learn and adopt modern DevOps tools and integrate them into their day-to-day activities. The landscape of options is broad, but the technology stack discussed here can provide a valuable toolkit to help DBAs increase their efficiency and productivity both locally and in the cloud. ■

CRAIG S. MULLINS

Craig S. Mullins is president of Mullins Consulting, Inc. He's an IBM Gold Consultant and the author of two best-selling books, *DB2 Developer's Guide* and *Database Administration: The Complete Guide to DBA Practices & Procedures*. Website: www.mullinsconsulting.com

Managing Database Performance

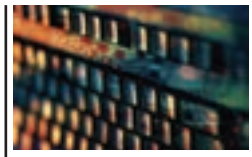

DBA CORNER

MANAGING THE PERFORMANCE of database systems and applications is a significant job responsibility for DBAs. From a database perspective, there are three basic performance components that must be performed:

1. Monitoring the database management system and the applications accessing it to find problems as they arise. This is typically referred to as performance monitoring.
2. Analyzing performance data (logs, trace records, reports, etc.) from the system to determine the root cause of problems.
3. Assembling a corrective action to implement a fix to the problems.

There are database performance software products that can aid with all three of these components. But you must be careful to fully understand the capabilities of any database performance management solution, as some simply monitor, others just analyze data or provide fixes for problems, and others deliver functionality combining all of these tasks.

You can also break down database performance management software by the category of performance issues it addresses. Database performance problems can occur in any of the following three areas:

- The DBMS itself, which must interact with other system software and hardware, requiring proper configuration to ensure it functions accurately and performs satisfactorily. Additionally, there are many database system parameters used to configure the behavior of the DBMS and the resources it has available to it. This includes criteria such as memory capacity, I/O throughput, and locking of data pages.
- The database design and schema, including database parameters, table designs, and indexing, can all impact database performance. How the data is organized must also be managed; as data is modified in the database, its efficiency will degrade. Reorganization and defragmentation are required to periodically remedy disorganized data.
- Finally, the SQL and application code itself can cause performance issues. Coding efficient SQL statements can be complicated because there are many different ways to write SQL that return the same results. But the efficiency and performance of each formulation can vary significantly. DBAs need tools that can monitor the SQL code that's being run, show the access paths it uses, and provide guidance on how to improve the code.

Database performance tools can identify bottlenecks and points of contention, monitor workload and throughput, review SQL performance and optimization, monitor storage space and fragmentation, and view and manage your system and DBMS resource usage. Of course, a single tool is unlikely to perform all of these tasks, so you likely will need multiple tools (perhaps integrated into a functional suite) to perform all of your required database performance management tasks.

Without proactive tools that can identify problems as they occur, database performance problems are most commonly brought to the attention of the DBA by end users. The phone rings and the DBA hears a complaint that is usually vague and a bit difficult to interpret, such as, "My system is slow today" or, "My screen isn't as fast as it used to be." To resolve such issues, DBAs need tools that can help uncover the exact problem and identify a solution. Database performance management tools can be used to find the root cause of such problems as well as to deploy a solution to fix the problem.

Furthermore, many organizations use multiple DBMS products in production, and the same DBA team (and sometimes even the same exact DBA) will have to ensure the performance of more than one DBMS (such as Oracle and SQL Server or Db2 and PostgreSQL). But each DBMS has different interfaces, parameters, and settings that affect how it performs. Database performance tools can mitigate this complexity with intelligent interfaces that make disparate components and settings look and feel similar from DBMS to DBMS.

There are many providers of database performance management tools, including the DBMS vendors (such as IBM, Microsoft, and Oracle), large ISVs (such as BMC and CA), and a wide array of niche vendors that focus on DBA and database performance software (for example, Quest, IDERA, and Navicat).

The exact database performance management solutions you should use depend upon the database systems you utilize, the size of your organization, the amount of data managed, your service level agreements, and your budget. But managing production databases without performance tools is a recipe for failure. ■

TODD SCHRAML

Todd Schraml has more than 20 years of IT management, project development, business analysis, and database design experience across many industries from telecommunications to healthcare. He can be reached at TWSchraml@gmail.com.

Beware the Frankenmart!

Novice developers are often confused about database design. And when they mature into experienced personnel, they unknowingly pass on bad habits to the next generation of developers. The problem at hand is a very human one. When putting together data marts, particularly data marts intended to be dimensional, some developers tend to think the way a typical end user might think. End users are very creative. If end users have a need, a function, a task at hand that the given application does not directly support, they will find a way to instantiate that missing functionality. It may require very unexpected uses of the existing solution, but the need will win out somehow. Developers may or may not understand the tenets and approaches that comprise a "dimensional data mart"; and often those data marts are implemented inside relational DBMSs. Therefore, the developer often feels that anything the DBMS will do is "fair game" for their data mart.

A relational DBMS is a very generic tool, whereas a multidimensional data mart is a specific kind of approach. Therefore, including "anything" that is relational actually spans a great many features that are not multidimensional in nature. The result of such a composite approach likely results in a data mart that has dimensional nomenclature applied to it, but that in practice is really an amalgamation of third normal form, denormalized, and dimensional structures. Most optimistically, one could think of these data marts as a pastiche, a bricolage, or to take a little poetic license, a data-based portmanteau morph, containing the best of all possible elements needed to deliver a solution. In practice, results are often less lofty. The data mart may be a complete pile of junk, or a Frankenstein's monster comprised of pieces from here and there that function, sort of. These Frankenmarts exist in many places. The limitations of these beasts often are associated with two main concepts. One, writing SQL to access data correctly means having a detailed understanding of all the pieces. Consequently, there may be a lengthy learning curve for a new querier to become functional. And next, because of the numerous differing approaches contained within the single solution, scaling may not be straightforward. The simplest of changes to the solution's requirements might force massive, or even complete, refactoring that could prove quite costly.

Unlike our end user, data mart builders must understand what they are working to accomplish. The DBMS is not going to magically guide them to a solution. The builder is responsible for knowing how dimensional techniques work, why they work, and what options may exist within the dimensional framework. The DBMS has a lot of functions for serving all sorts of purposes. Architecture means in part that there is a place for the elements used, and the elements used are all in their proper place. The person serving as the data architect must be the enforcer of the rules, or else there will be none. A Frankenmart has pieces put together, but likely not in their proper place, and that is the ultimate weakness. Simple is almost always best. If the target is a multidimensional data mart, then work toward exactly that. Exceptions may arise, but those exceptions should be very rare and should be very well-documented. Therefore, the best results are most often simple star schemas, with no snowflaking and no strange normalized components dancing around facts. There should be no expectation of direct fact table-to-fact table joins. And ideally, even though they are "legal moves," there should be an absolute minimum of bridge and outrigger tables. Enforcement is completely in the hands of the data architect in charge, and these individuals need to have the fortitude and belief to imprint their will on the solution's database design. ■

Early bird registration ends January 25

Online agenda now available

SHARE Phoenix 2019

March 10-15, 2019

PHOENIX CONVENTION CENTER

Learn how to increase efficiency, agility and cost savings through 500+ technical sessions focused on enterprise IT's hottest topics such as Pervasive Encryption, Cloud Technology in the Enterprise, Data Privacy, API Economy, Zowe and Open Source, as well as IBM Z topics.

LEARN MORE AND REGISTER AT **event.share.org**