PAPER

Redmond MAGAZINE

ZEN AND THE ART OF EMBEDDABLE DATA

Multiplatform, embeddable database for self-managed, small-footprint environments

By John K. Waters

EXECUTIVE SUMMARY

Actian Zen is a multiplatform, embeddable database for self-managed, small-footprint environments that was developed over more than a decade from a group of technologies created and acquired by Actian, a Silicon Valley-based maker of data management, analytics, and integration products. It was rebranded in 2017 to underscore its essential relationship with developers facing the unique challenges of creating applications for the Internet of Things (IoT) and Edge Computing.

This is a solution for a new ecosystem, but one with deep roots in technologies that have been around for more than 30 years. Because of those roots, Actian Zen is unique among embeddable database solutions in several ways. This white paper looks at the origins of this product, the capabilities it currently provides, and the environment into which it has emerged.

DEFINITIONS

Almost inevitably, some confusion arises when the discussion turns to edge of the network, and it's worth taking a moment to define a few terms.

Embedded

The terms embedded database and embeddable database refer to a database and/or a management system (DBMS) that is integrated into, or very tightly with, the application that needs direct or fast access to the data it contains. The DBMS is therefore dedicated to the app using it.

It's important to remember that an embedded database can be an in-memory database, reside on solid state or traditional disk storage, or be some combination of both.

Also, an embedded database is a critical part of an intelligent embedded system, where special-purpose software with a dedicated function is installed on a piece of hardware.

ACTIAN ZEN IS A MULTIPLATFORM, EMBEDDABLE DATABASE FOR SELF-MANAGED, SMALL-FOOTPRINT ENVIRONMENTS THAT WAS DEVELOPED OVER MORE THAN A DECADE FROM A GROUP OF TECHNOLOGIES.

Cover Image: seewhatmitchsee / Shutterstock.com

IoT

The Internet of Things is part of the popular lexicon now, but just to be clear, IoT refers to the vast and ever-expanding network of physical devices linked through the Intranet or Internet. There are billions of them now, everything from smart phones to temperature sensors, appliances to cars. The software installed on them gathers, stores, and analyzes data, either locally or through a connection with separate compute platforms and associated databases.

EDGE COMPUTING PUSHES THE INTELLIGENCE, PROCESSING POWER, AND COMMUNICATION CAPABILITIES OF AN EDGE GATEWAY OR APPLIANCE DIRECTLY INTO DEVICES.

One subcategory worth keeping in mind is the IoT gateway (sometimes called an intelligent gateway or a control tier). It can take the form of a physical hardware appliance or a piece of dedicated software. It serves as the connection between the cloud and the universe of controllers, sensors, and/or intelligent devices.

Edge Computing

Edge Computing refers to a distributed information architecture in which the processing of data occurs where it is generated, far from a centralized datacenter, at the "edge" of the network. That processing is done on IoT devices and remote gateways. Edge Computing has also been described as a decentralized data persistence that occurs on or near the devices that generate the data.

Fog Computing, a term coined by Cisco Systems in 2014, effectively means the same thing as Edge Computing, though it's more directly associated with Cloud Computing and has a Cloud-side view of where in the Edge intelligence will preside. As Cisco defines it, the Fog "extends the cloud to be closer to the things that produce and act on IoT data." The term has been called a marketing take on Edge, but it has been catching on. ("Fogging" has even become a verb.)

To be clear, Fog Computing pushes intelligence down to the local area network (LAN) level of the network architecture, and the data is processed through a "fog node" or IoT gateway; Edge Computing pushes the intelligence, processing power, and communication capabilities of an edge gateway or appliance directly into devices. More than likely, there will be an ebb and flow of intelligence between gateways and devices such that the distinction between Fog and Edge terms are largely semantics.

IOT AND COMPUTING AT THE EDGE OF THE NETWORK

Small-footprint embeddable databases are not new, of course. During the 1990s and early 2000s, embedded DBMSs were used to manage data flows over wired networks for a range of media-intensive applications. But today, virtually all transactions among IoT devices, media-intensive and otherwise, require this level of data access.

It could be argued that the primary advantage of an embedded DBMS is reduced network latency and faster application performance. If the application and the data it needs to access reside on separate machines, the app will be slow; if they're tightly integrated, the app will be fast. Speed matters, especially when it comes to IoT devices. Lagging response times from fitness bands and smart thermostats are annoyances among users of those products, but it's downright dangerous in autonomous vehicles and medical devices.

But fast data access isn't the only advantage of an embedded database. An embedded DB comes with a small footprint, which is critical in resource constrained environments. Also, it's distributed with the application, so there's no need for the end user to install it separately—in fact, the database is invisible to the end user. And because it's linked directly to the application as an in-process software library, it requires almost no database administration. Finally, because the database is dedicated to the application, there is less chance of security vulnerabilities being introduced by other applications, as would be the case with a communal database typically used in enterprise data warehouses, data lakes and other multi-purpose environments.

THE ROOTS OF ZEN

Actian was the original commercial supporter of Ingres, an open-source, SQL relational DBMS developed in the 1980s at the University of California, Berkeley, for large commercial and government applications. (There's still an active Ingres user/developer community that Actian supports.) In the early 2000s, the company shifted its focus to the emerging opportunities and challenges of the Big Data phenomenon, and then to IoT and Edge Computing. The list of its acquisitions over the past decade includes ParAccel, a provider of a DBMS for Big Data analytics; VectorWise, an Ingres-based SQL relational DBMS aimed at analytical database applications; and Pervasive Software, a vendor of cloud-based and on-premises software for data management and analysis and creator of the DBMS known as PSQL.

PSQL is optimized for embedding in applications, and it is the core of Actian Zen. Its architecture comprises a transactional database engine and a relational database engine, both of which can access the same data, but through different methods. The transactional engine, known as the MicroKernel Database Engine (MKDE), interacts directly with the data and does not require fixed data schema to access it. The relational component, known as

ZEN HAS EMERGED AT THE EDGE OF THE NETWORK AS A FOCUSED GROUP OF EMBEDDABLE DATABASE MANAGEMENT SOLUTIONS WITH AN IMPRESSIVE RANGE OF CAPABILITIES THAT SHOULD MAKE LIFE EASIER FOR DEVELOPERS.

the SQL Relational Database Engine (SRDE), is a relational DB accessed via Structured Query Language (SQL) queries. Those queries are parsed and sent to the MKDE to run.

Early versions of the product were called Btrieve and then Pervasive PSQL, and the company still provides support for products purchased under the Pervasive name. With Zen, the company has extended the essential features of the embedded editions of "classic PSQL" to a new generation of intelligent apps in the Cloud and pared down to support smart devices at the Edge. The new name is actually an acronym meant to reflect the product's key features for IoT and Edge Computing: Zero-DBA, Embeddable, and Nano-footprint.

"Zen" is now the rubric for a product family that includes several editions of the embeddable database: Zen Edge Server for IoT; Zen Enterprise Server for Windows, Linux, and Mac OS; Zen Core for Android and iOS (Zen Core can be adapted to virtually any device running a POSIX compliant variant of Linux); Zen Workgroup for Windows Server; and the Zen Cloud Server (on VMware vSphere, Microsoft Hyper-V, and Citrix XenServer). The company markets these products to system integrators (SIs), independent software vendors (ISVs), original equipment manufacturers (OEMs), and enterprise customers building in-house applications that require an embedded DBMS.

ZEN AT THE EDGE

Zen has emerged at the edge of the network as a focused group of embeddable database management solutions with an impressive range of capabilities that should make life easier for developers facing the vicissitudes of an evolving mobile application front end and IoT ecosystem.

Common data type

One clear breakthrough feature is Zen's ability to deliver a common data type and file format across all platforms, as the company puts it, "from the edge to the gateway to the enterprise." This capability eliminates the typical ETL overhead involved when moving data among operating environments or versions of the database. Data moved from, say, an ARM device (such as a Raspberry Pi) to a Linux server or virtual machine in a Public or Private Cloud requires no ETL. Again, as the company puts it, "just copy the data and go."

SQL/NoSQL

Another impressive innovation is Zen's NoSQL programmatic API-based database access and SQL relational database access to a single data set. Applications can be designed to take advantage of either or both access methods for database transactions, reporting, and data analytics. It's a kind of one-two performance punch that provides both speed and flexibility. The NoSQL support includes several Software Development Kits (SDKs), which provide developers with direct access to the data without requiring them to go through a relational layer. The result is fast read and quick insert, update and, delete performance alongside full ACID response on writes and ANSI SQL queries. SQL access is provided via ODBC and JDBC and NoSQL access via the Btrieve and Btrieve 2 APIs.

Modularity

Modularity is another noteworthy attribute of the Zen products. All of these solutions are built with a single, modular architecture that scales from a core set of libraries capable of single-user client data management to a direct key-value store engine, and even up to a full-edged, enterprise-grade server capable of supporting thousands of users on multi-core, VM cloud environments (such as VMware vSphere, Microsoft Hyper-V, and Citrix XenServer).

Small Footprint

When the conversation turns to embeddable databases, much is made of a product's footprint, and rightly so; these databases are joining applications in extremely tight quarters. The general descriptions range from small to "nano." There's no real standard here, but Zen's range from under 2MB, available on Android and iOS, to 175MB, a minimum for the Enterprise edition.

Other Features Worth Noting

- **Zero Database Administration:** The Zen DBMSs are designed for environments where users can't be expected to manage a database. Zen runs without end-user interaction, notice, or DBA supervision.
- **Security:** Zen comes with a full complement of security, encryption, management, and monitoring tools,
- **Defragmentation without down time:** Files can be defragmented while the system remains online. Reads and writes can operate continuously during the defrag process.
- **Backward Compatibility:** The Zen DBMSs allow ISVs to take advantage of new hardware architectures, OS platforms, and computing environments (64-bit, multi-core, VM, Cloud) without application changes. Moves to new platforms can occur with no new application code.
- **Customized Installation:** Many ISVs building applications for small business users, or even for large enterprises with sprawling branch and remote field environments void of any local IT expertise, want to hide the complexity of a database from those users. An Actian Zen installation can be hidden in the ISV application (it's called a "silent install"), and the developer can select the

components needed to eliminate unneeded features or keep the footprint to a minimum.

- **Multiplatform Support:** The Zen products support a number of hardware and software platforms that are critical to developers of IoT and Edge Computing applications, including Windows, Linux, macOS, Android, iOS, Raspbian, VMware, ARM, and Intel.
- **ACID-Compliant:** The Zen embedded databases conform to the standard for reliable database transactions, which comprises four properties: Atomicity, Consistency, Isolation, and Durability. ACID compliance means the information collected is guaranteed to be accurate.
- **Client Reporting Engine:** This feature allows customers to run SQL queries against client cache and offload large reporting workloads to a separate machine, reducing the impact on the main server engine. Once it is set up, the Client Reporting Engine presents to SQL clients the same databases as the main server, so all SQL queries can be run on it.
- **Enterprise features** like data replication, data auditing and live database backups are also available with server editions.

CONCLUSION

Actian has done an excellent job of leveraging the PSQL heritage of its Zen embeddable database to produce a suite of products with practical features and capabilities that will solve a myriad of problems for developers of IoT and Edge Computing next-generation applications. The company's decision to tailor its original product to the specific requirements of popular platforms and use cases was a smart one that's sure to get the attention of ISVs, OEMs, and enterprise developers, who are heading for the edge of the network at warp speed. Perhaps more importantly, the company has created a solution that is flexible enough to evolve with the market the company seeks to serve.

Find out more:
https://www.actian.com/data-management/zen-embedded-database/

John K. Waters is a freelance journalist and author based in Silicon Valley. He is editor-at-large for Application Development Trends, and a contributing editor to Campus Technology, T.H.E. Journal, and Law Technology News. He covers a wide range of topics in information technology and is the author of more than a dozen books. He writes news and features about software development trends, application and network security, infrastructure and database technologies, cloud computing, open source software, standards and governance, virtualization, and the people working in, and issues affecting, the information technology business.