iguazio

ACCELERATE BUSINESS VALUE WITH SERVERLESS COMPUTING

Written by Jason Bloomberg, President, Intellyx

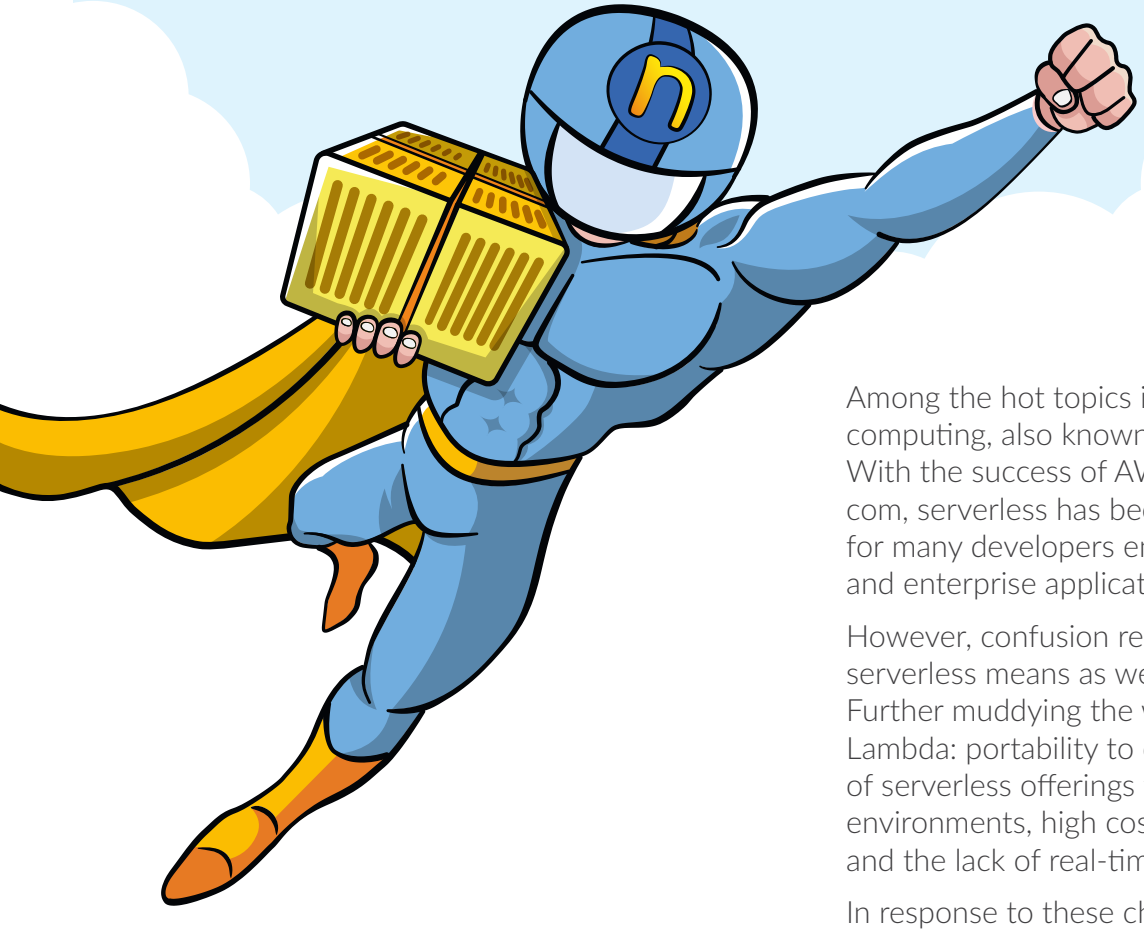

Among the hot topics in IT today is serverless computing, also known as Functions-as-a-Service. With the success of AWS Lambda from Amazon.com, serverless has become an essential must-have for many developers engaged in a range of web-scale and enterprise application initiatives.

However, confusion remains over exactly what serverless means as well as its pros and cons. Further muddying the waters are the limitations of Lambda: portability to other cloud providers, a lack of serverless offerings for edge and on-premises environments, high costs, operational challenges, and the lack of real-time analytics.

In response to these challenges, open source offerings like Nuclio from Iguazio seek to resolve the limitations of public cloud-based serverless computing while also offering unprecedented performance. The end goal is to reduce barriers to serverless and bring it to the forefront of high-performance, real-time computing.

As technology continues its inexorable, Moore's Law-driven exponential progress, enterprises are increasingly able to leverage vast technical resources to drive a new class of digital application.

Cloud computing, Agile and DevOps development approaches, and low-code technologies are all part of the fabric of available technical resources and expertise that drive such applications.

Customers, employees, and everyone else interacting with the organization now expect applications to operate in real-time and scale without limit. Such applications span on-premises, cloud, and edge computing resources, as hybrid IT drives workload-centric computing across the enterprise.

And yet, the promise of modern technology still lags behind the reality for most organizations. Traditional application infrastructures remain slow to build, and developers must spend much of their time on infrastructure issues like scaling, input/output, and other activities that add little to no value to the business.

For software engineers to step up to the plate in their organizations, they should focus instead on value-generating business logic. In contrast, all infrastructure tasks are little more than plumbing that the underlying platform should handle automatically.

Such is the motivation behind one of the hottest new areas in cloud computing: serverless computing. With serverless computing, developers simply create and access individual software functions – all implementation details are hidden from view.

The name serverless, however, is actually a misnomer, as there are still servers behind the scenes.

Unlike IaaS, however, when using serverless computing, developers don't have to instantiate server instances, or worry about VMs at all. Unlike PaaS, furthermore, there is no platform churning away in the background running up the bill.

With serverless – also known as 'Functions-as-a-Service' (FaaS) – the cloud infrastructure fires up the necessary technology at the time and place the user needs it, and then turns it off when they're done.

Serverless computing addresses limitations of modern software development by moving communications logic, scalability, parallelization, and other infrastructure tasks – the plumbing – out of the program and into the platform. It also supports common, popular languages like Go, Python, and NodeJS.

Today, the primary source of serverless capabilities is Amazon Lambda. Amazon deserves credit for bringing serverless computing to market, establishing it as a viable approach. However, as with many early entrants to new markets, it has weaknesses as well as strengths.

Fortunately, there is an alternative to public cloud FaaS: Nuclio.

NUCLIO: THE OPEN SOURCE ALTERNATIVE TO PUBLIC CLOUD SERVERLESS

While FaaS point solutions are becoming increasingly common, the number of true, 'bare metal' (i.e., install wherever you like) serverless platforms is small. This paper will focus on one open source solution in particular: Nuclio.

Nuclio is the brainchild of Iguazio, the creator of a high-performance data services platform for real-time applications. However, while Iguazio's data platform is a commercial product, the company chose to release Nuclio as open source, thus welcoming input from a broad community of contributors.

As a true serverless computing platform (as opposed to a FaaS point solution), Nuclio is cloud neutral. In other words, an engineer can install it either on-premises or in their choice of public or private clouds, including within various hybrid, multi-cloud and edge scenarios.

In fact, Nuclio's role as a platform addresses the 'if it's serverless, then why do you need servers' conundrum. Even though Nuclio allows for bare metal installation (thus requiring servers), it provides all the benefits of the abstractions inherent in public cloud serverless offerings, including scalability, simplicity, and developer productivity.

Just as importantly, Nuclio addresses the drawbacks of public cloud serverless. As an open source offering, Nuclio has no vendor lock-in. In addition, every element of the platform is available for monitoring and management, because there's no cloud provider intentionally obscuring its inner workings.

Nuclio can also run on a private network if necessary, offering a reduced threat surface as compared to publicly available functions. Finally, using Nuclio in conjunction with the Iguazio data platform resolves all the issues of state management and persistence inherent to public cloud serverless offerings.

COMPARING NUCLIO TO LAMBDA: THE PROS OF SERVERLESS COMPUTING IN THE CLOUD

Serverless computing is born of the public cloud. AWS is leading the way with its Lambda offering, with Microsoft, Google, and IBM not far behind. Here are its most important benefits:

CLOUD STRENGTHS

As you would expect, such cloud-based offerings take advantage of all the strengths of the public cloud: massive scalability, transparent elasticity, pay-as-you-go pricing, and straightforward, API-driven access to all the other services each cloud has to offer.

Serverless computing then builds on these advantages. Serverless architectures offer automatic scale up and down in response to the current load on a function-by-function basis – where the cloud provider only charges for the handful of milliseconds of compute time each function uses when it is running.

COST

The operational context of serverless computing is entirely hidden from view. In essence, each function is itself a live, production instance – fully able to participate in full-blown applications, while only running (and thus running up the bill) when such applications actually need them to run.

As a result, there are two ways that serverless computing saves money over other cloud approaches. First, it essentially turns off when not in use. But even more importantly, it requires no time and effort on the part of your operations personnel.

In essence, the cloud provider, rather than your personnel, is responsible for setting up all the runtime infrastructure and ensuring the functions always scale to meet the demand.

PRODUCTIVITY

Serverless computing also has an important top-line benefit: increased developer productivity. Functions are simple to incorporate into applications, and natively support microservices architectures.

Serverless computing empowers developers to focus on the value-added part of their jobs, without having to worry about the plumbing.

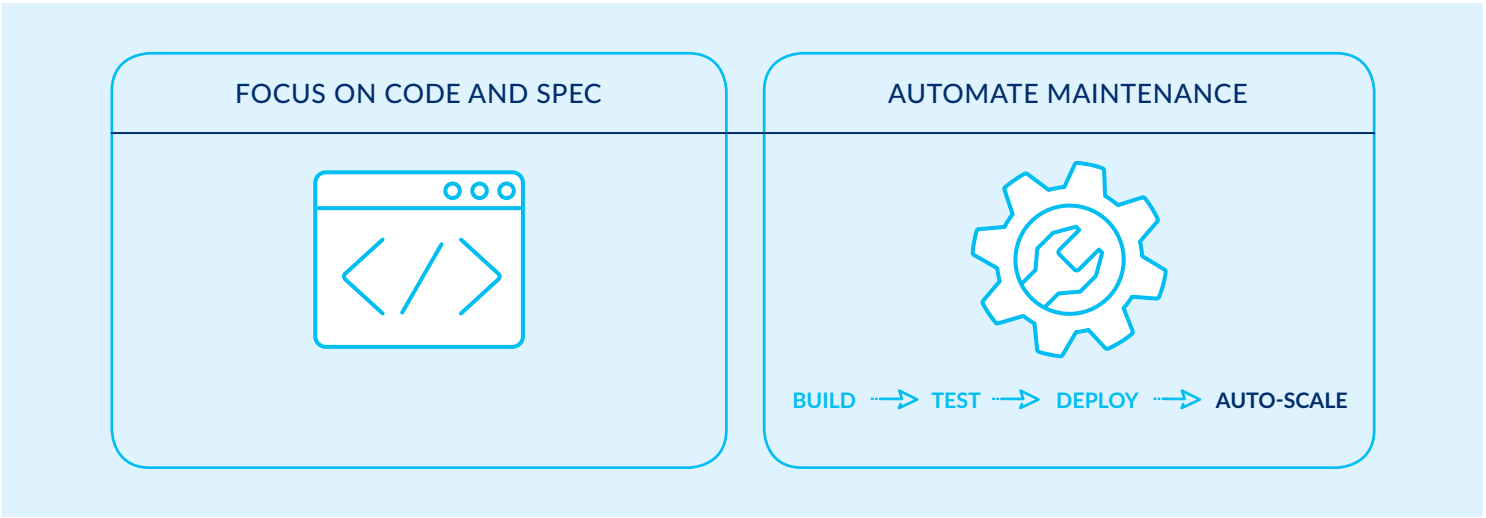

THE CONS OF CLOUD-BASED SERVERLESS COMPUTING

When relying upon a public cloud to deliver serverless computing, however, there are also some important disadvantages to consider. After all, the public cloud isn't suitable for every company or every use-case – and furthermore, many organizations are wary of putting all their eggs in the Amazon basket.

PERFORMANCE

Even though the cloud offers scalability and elasticity, serverless computing still suffers from significant performance concerns.

The reason is simple: the infrastructure must spin up each function as needed, an action that can take a few milliseconds – a few milliseconds which can potentially slow everything down.

Sometimes, in fact, this 'spin up' can take even longer. For example, when the underlying runtime (a Java runtime, for example) must spin up before the runtime can launch the function, which can take several seconds. Furthermore, whether the runtime is running or not is hidden from view, making it difficult to plan for.

Even if the delay be only a few milliseconds, it can contribute to unacceptable latency, especially in the case of real-time applications. A few milliseconds of latency in a video feed, for example, can lead to the dreaded spinning wheel. In the more general case, streaming data applications often cannot tolerate such unpredictable latency.

COST

As with other cloud services, serverless computing does not necessarily provide the lowest-cost option, especially at the high-performance end of the spectrum. Simply put, pay-as-you-go isn't best for everyone.

For certain workloads, bulk provisioning of VMs will be the more cost-effective option – and at the extreme, physical servers may even be the least expensive choice.

Serverless functions are essentially an 'a la carte' option – but sometimes the full meal is a better deal.

MANAGING STATE AND PERSISTENCE IN GENERAL

The next challenge: managing state information. Serverless functions are inherently stateless, which means that they don't store information internally, and thus they cannot keep track of where a particular function call fits into a process that might string together such calls.

From the perspective of the public cloud providers, such statelessness isn't a problem – they simply offer various databases or other data persistence mechanisms to which engineers can connect their serverless functions.

This build-it-out-of-Legos approach works well in many situations, but further adds to the complexity of any serverless deployment, and furthermore, increases cloud provider lock-in – a significant concern for many IT executives.

SECURITY

Security may also be a concern for cloud-based serverless computing – not because the cloud provider's own security isn't up to snuff, but simply because serverless computing expands the enterprise threat surface.

In essence, every connection between a company and its IT resources outside the company presents a possible point of vulnerability that the security team must protect. Serverless functions thus act as 'pinpricks' in the corporate threat surface – small but potentially numerous and difficult to wrangle.

MANAGEABILITY

Finally, the fact that cloud providers hide the inner workings of serverless computing presents a problem for monitoring and management. Cloud providers are telling their customers 'trust us' – but responsible organizations should follow the mantra 'trust but verify.'

The 'verify' part of that equation is particularly challenging in a cloud-based serverless environment.

Nuclio self-service
portal and marketplace

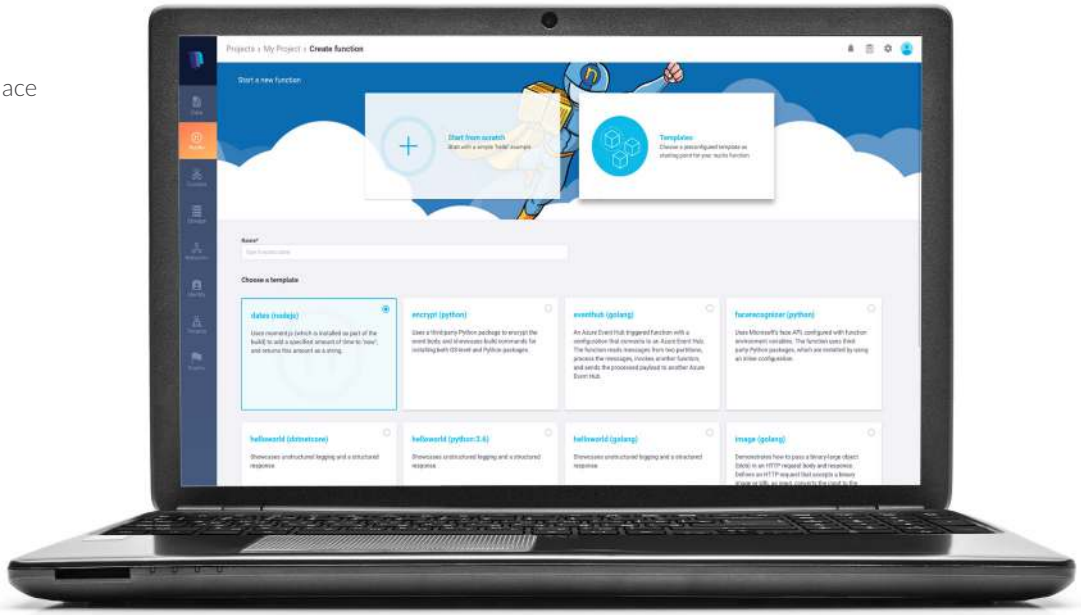

‘BARE METAL’ SERVERLESS COMPUTING: THE BEST OF BOTH WORLDS?

If you review the cons of serverless computing, all but one are directly due to the fact that serverless functions run in a public cloud – the exception being the performance hit due to spinning up the execution environments for the functions.

To address these issues, as well as to give organizations a wider range of serverless deployment options, a number of open source and commercial projects have sprung up to enable them to install and configure the entire FaaS platform wherever they like – either on-premises or in their cloud of choice.

This pattern in which the market addresses vendor lock-in concerns by rolling out open source alternatives is a familiar one, particularly in the enterprise space. After all, Linux itself followed this pattern.

As this scrappy open source operating system matured, it resolved issues with various Unix flavors, Windows NT, and older server OS’s to the point that most vendors eventually chose to just adopt Linux and fight their battles elsewhere.

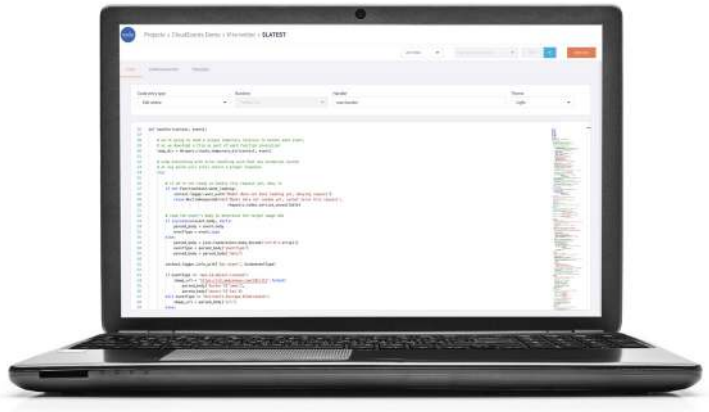

NUCLIO: ARCHITECTED FOR PERFORMANCE

Today, however, we're in the earliest days of open source serverless computing. In fact, many of the projects on the workbench today (both open source and proprietary) are FaaS, but not true serverless in the sense that they don't offer a managed platform that fully abstracts both VMs as well as the function runtimes.

Nuclio fills this gap. In fact, the most significant benefit of Nuclio over public cloud serverless services – as well as other open source serverless projects – is how it resolves the performance bottleneck.

Whereas other serverless platforms (including those in public clouds) operate on a 'start up upon function call' principle, Nuclio runs on a 'real-time processing' basis. As a result, Nuclio is far faster than alternatives – up to one hundred times faster.

The diagram below illustrates the function processors and platform services at the core of the Nuclio architecture.

In Nuclio, all functions are event-driven. They respond to event triggers, data messages, or records from any of several different types of event sources. The function engine in the Nuclio processor then processes such events in real-time.

PORTABILITY ACROSS ENVIRONMENTS

Furthermore, all of this power and speed can run as a standalone binary, either for true bare metal installs or for deployments in specialized environments like IoT gateways.

Nuclio can also run as a package within a Docker container, or integrated with a container orchestrator like Kubernetes. In fact, Nuclio's support for containers and orchestrated container environments makes its functions seamlessly compatible with microservices, which typically run within containers.

Because microservices are, by definition, parsimonious (as small as possible but no smaller), serverless functions are an ideal tool for building streamlined, simple microservices.

Given the performance, control, and manageability of Nuclio, the combination of Nuclio with Kubernetes or other container orchestrators yields perhaps the best combination of capabilities for building high-performance microservice-based applications available on the market today.

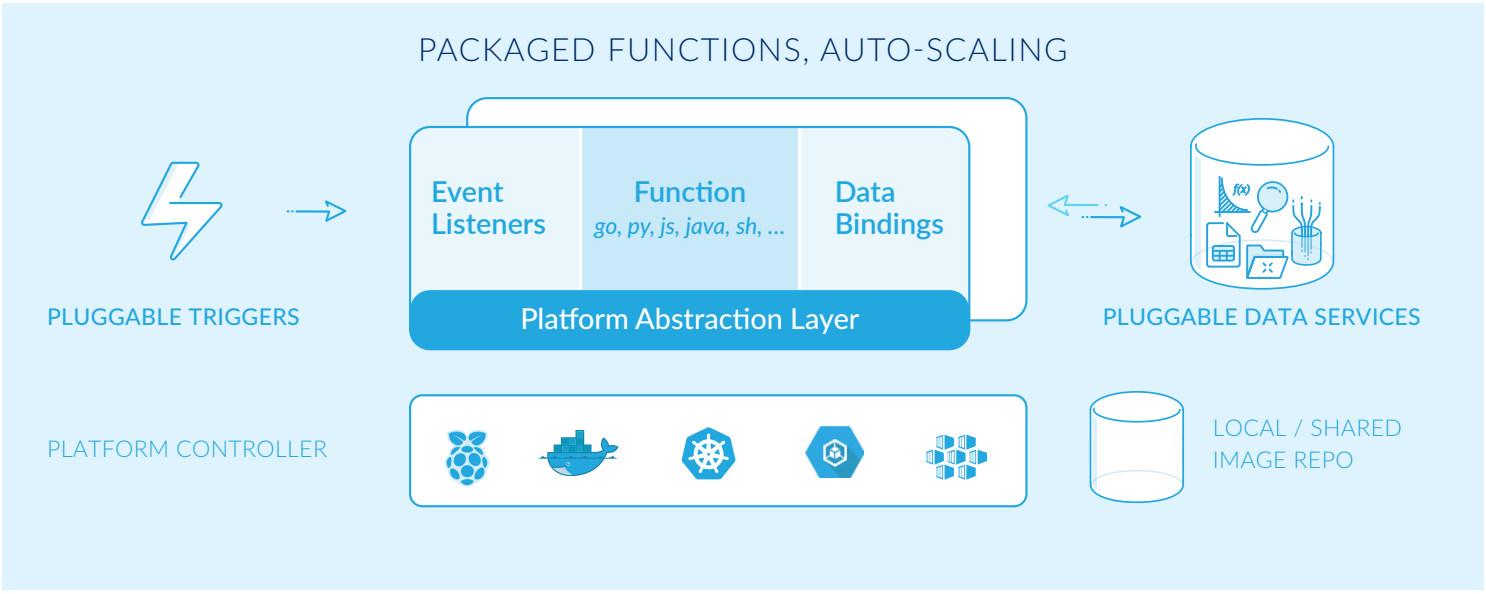

The core of Nuclio's architecture (source: Iguazio)

THE INTELLYX TAKE: THE FUTURE OF REAL-TIME DATA-CENTRIC APPS

While still relatively new on the market, serverless computing has proven its value via a combination of scalability, productivity and cost benefits. Nuclio in particular promises to address many of the limitations of AWS Lambda, further cementing serverless computing's place in the enterprise.

The open source Nuclio community is also shaping up. While Iguazio will continue to have a hand in the open-source Nuclio effort, the community around the platform will continue to grow and diversify, making it a stronger offering by the day.

Meanwhile, Iguazio will remain leading Nuclio experts, and we can also expect to see a commercial offering from Iguazio that leverages Nuclio at its core.

This commercial version of Nuclio promises to have even greater power and enterprise-class capabilities than the core Nuclio platform, because Iguazio's real-time data platform neatly complements Nuclio's capabilities and adds the necessary state management and general persistence necessary for serverless computing to support enterprise-class applications.

Secondly, we can expect to see additional security and compliance capabilities from Iguazio that will add enterprise-level functionality to Nuclio that larger organizations typically require.

Regardless of the specifics of the Iguazio roadmap for Nuclio, however, Nuclio stands well on its own as an open source project. It is already mature enough to take its place with Kafka, Spark, Hadoop, and other high-performance, enterprise-class open source initiatives.

Jason Bloomberg is a leading IT industry analyst, Forbes contributor, keynote speaker, and globally recognized expert on multiple disruptive trends in enterprise technology and digital transformation. He is founder and president of Agile Digital Transformation analyst firm Intellyx. He is ranked #5 on Onalytica's list of top Digital Transformation influencers for 2018 and #15 on Jax's list of top DevOps influencers for 2017.
Mr. Bloomberg is the author or coauthor of four books, including The Agile Architecture Revolution (Wiley, 2013). His next book, Agile Digital Transformation, is due within the next year.

ADDITIONAL RESOURCES

Nuclio on Github https://github.com/nuclio/nuclio

Serverless whitepaper by the Cloud Native Computing Foundation https://github.com/cncf/wg-serverless/tree/master/whitepaper

ABOUT IGUAZIO

The Iguazio Continuous Data Platform digitally transforms businesses by analyzing data from a variety of sources and types to create actionable insights. Through its platform, Iguazio simplifies the development and deployment of high-volume, real-time, data-driven applications to extend the cloud experience to the edge and on-premises environments. Iguazio accelerates the digital transformation of manufacturing, healthcare, financial services and others by supporting rapid development of applications for the Internet of Things, autonomous systems, media, cybersecurity and smart mobility. The company was founded in 2014 and has offices in the USA, Singapore and Israel.