



At Ignite, Microsoft convincingly addresses the multicloud imperative by, James Kobielus

November 7th, 2019

Microsoft Corp. has fresh momentum in the cloud wars. For starters, Microsoft was recently awarded the U.S. Department of Defense's Joint Enterprise Defense Infrastructure or JEDI contract. For another, the company has just announced a broad range of enhancements to its Azure cloud portfolio that will almost certainly help to win it new business in the tooth-and-nail [...]

At Ignite, Microsoft convincingly addresses the multicloud imperative

Microsoft Corp. has fresh momentum in the cloud wars.

For starters, Microsoft was [recently awarded](#) the U.S. Department of Defense's Joint Enterprise Defense Infrastructure or JEDI contract. For another, the company has just announced a broad range of enhancements to its Azure cloud portfolio that will almost certainly help to win it new business in the tooth-and-nail battle with Amazon Web Services, Google Cloud Platform and others.

This week at its [Ignite 2019 conference](#), Microsoft launched both new and enhanced data-centric solutions for its cloud solution portfolio. For information technology professionals managing their firms' investments in cloud technologies, these were the chief announcements of interest at Ignite 2019:

- **Deepening its hybrid cloud functionality:** Microsoft strengthened its hybrid cloud solution, the Azure Stack Hub. It announced new [Azure Stack Hub services](#) for cloud database management, backup and recovery, real-time data ingestion, stream analytics, Kubernetes and virtual desktops. As AWS gets closer to making its Outposts hybrid-cloud offering generally available, it will face an uphill battle trying to match Microsoft's already mature rival offering.
- **Expanding its multicloud reach:** Microsoft took a bold step into supporting multicloud deployments of its Azure services. Currently in preview, the new [Azure Arc](#) supports cloud-to-edge reach for Azure services — starting with Kubernetes and Azure SQL Analytics, without sacrificing unified management, governance, security, compliance, auditing and control of these services wherever they live. The announcement of this offering in preview propels Microsoft directly to the forefront of the emerging enterprise multicloud. Once it's in production, Arc will allow customers to leverage their investment in Azure services by, to use Azure advocate Jeramiah Dooley's term on theCUBE, "cherry-picking" which of them to deploy them to third-party public clouds, which to heterogeneous Kubernetes clusters on premises and in the cloud, and which to disparate edge infrastructure. As long as Arc comes to market within the coming year, it will be able to win business for Microsoft against such [first-to-market multicloud management providers](#) as IBM Corp. and Cisco Systems Inc.
- **Bringing its cloud services to the edge:** Microsoft took an important step in extending its Azure cloud services to the edge. It announced the new [Azure Stack Edge](#), a cloud-managed device offered on a hardware-as-a-service basis. Unlike VMware's still-undelivered "[Project Dimension](#)," this edge-computing offering is optimized from the start for AI. Each Azure Stack Edge node has network data transfer capabilities, incorporates a built-in FPGA for accelerated AI-inferencing and functions as a storage gateway. Microsoft is rolling out both a commercial series and rugged series for this new edge device with FPGA and GPU form factors for edge-based machine-learning inferencing. It offers automatic failover of the devices so that, in case of a loss of a server in an edge cluster, a VM can automatically move to another server. And it provides customers with the option of running applications on VMs or containers.
- **Scaling its virtual machines even further:** Recognizing that [virtual machines remain fundamental](#) to many cloud application environments, Microsoft gave Azure customers new reasons to stay with its virtual machine technology. This week, it made several announcements that improve the scalability, performance, efficiency, and deployment flexibility of its VM technology for Azure's infrastructure as a service. Most notably, Azure Generation 2 virtual machines enable larger workloads to run in VMs in Azure through support for larger memory, larger VMs (up to 12 terabytes), and provisioning of OS Disk sizes that exceed 2TB. Just as important in this era of the edge, Microsoft has announced a limited preview of a wide variety of virtual machine sizes tailored to meet diverse Azure customer needs, with its

At Ignite, Microsoft convincingly addresses the multicloud imperative

NV size family having been optimized for graphics processing unit-powered visualization workloads, such as CAD, gaming and simulation.

For developers and users, the chief announcements addressed multicloud-spanning data warehousing (the new [Azure Synapse Analytics](#), the evolution of Azure SQL Data Warehouse), automated data-science DevOps workflows (a.k.a. “MLOps,” a preview of new [Azure Machine Learning](#) enhancements), robotic process automation ([Power Automate](#), formerly called Microsoft Flow), and no-code development of AI-driven conversational user interfaces (the new preview of [Microsoft Virtual Agents](#)).

None of these were particularly innovative, considering that these solution spaces are well-established solution niches. But with these moves, Microsoft made it quite clear that it plans to keep evolving its cloud-to-edge development tooling in line with its back-end hybrid, multicloud, edge and virtualization platforms.

What’s especially important is the MLOps coordination and standard workflows within data science development teams. As Francesca Lazzeri, a machine learning scientist and cloud advocate at Microsoft, said on theCUBE, “This is something that you need to change within your company, within your organization. You have to always to make sure that data scientists, machine learning scientists are working closely with data experts, business experts, and everybody’s talking.”

Likewise, no-code approaches are essential for acceleration of AI and other cloud services application development throughout the enterprise. As Dona Sarkar, Microsoft’s advocate lead for Power Platform, said on theCUBE, “Chatbots are everywhere. The way people think about making them is, ‘Oh I have to go get Azure cognitive services and learn it deeply and become a AI expert and learn to like speak natural language processing stuff.’ But in fact, you can build a chatbot in five minutes using Microsoft Power Virtual Assist, which is fantastic and really cool. (With) AI Builder, I’m like, oh, I trained an AI model and deployed it and tested it on stage. That’s crazy and cool. And I learned to do that in five minutes and believe you me, I’m not a data scientist.”

For Microsoft’s customers, these announcements at this week’s Ignite enable them to unlock more value from their investments in Azure cloud services, especially for applications that hinge on big data analytics, machine learning and artificial intelligence. Just as important, the new Azure Arc mitigates the risk of lock-in to Azure, allowing them to move the business capabilities built in Microsoft tooling and deployed on Azure, to any third-party public cloud, any Kubernetes cluster or any on-premises platforms.

We recommend that Microsoft customers consider these announcements in the light of their own demand for hybrid cloud, multicloud and edge computing environments. With this week’s Ignite, Microsoft now has a credible platform and strategy for bringing computing truly anywhere and managing it all from end to end.

At Ignite, Microsoft convincingly addresses the multicloud imperative

Jim is Wikibon's Lead Analyst for Data Science, Deep Learning, and Application Development. Previously, Jim was IBM's data science evangelist. He managed IBM's thought leadership, social and influencer marketing programs targeted at developers of big data analytics, machine learning, and cognitive computing applications. Prior to his 5-year stint at IBM, Jim was an analyst at Forrester Research, Current Analysis, and the Burton Group. He is also a prolific blogger, a popular speaker, and a familiar face from his many appearances as an expert on theCUBE and at industry events."



James Kobiellus

@jameskobiellus

james.kobiellus@wikibon.org