Preparing for battle on the AI performance front?

Much of it has already been fought. Here's how – and what it means for IT leaders.

Ask an IT leader about their main concerns heading into the critical next phase of AI adoption and execution, and you'll hear many of the same themes. Data concerns tend to dominate the list, as in "which data?", "where should it come from?", and "what's the best way to blend so many different data streams?"

Strategy also looms large – leaders need to be confident in determining which use cases and outcomes to pursue, given the sprawling list of possible use cases.

They're also looking ahead to people-related issues: Do they have the talent they need? How will they get management support? How can they help ensure that a wider range of users benefit from new, embedded AI capabilities?

And then there's performance. Everyone knows that AI workloads require more processing power than core analytics tasks. You might expect that IT leaders are gearing up for a long, difficult period of integrating different hardware and software assets to ensure a high level of performance in AI environments. In reality, many are not – they simply assume that the cloud will address performance obstacles, especially if paired with a separate, AI-dedicated analytics appliance.

It's easy to see the appeal of this approach, and cloud will certainly play a valuable role in AI strategies. But it's also short-sighted, leading to a "shadow analytics" situation in which AI-related analytics capabilities are executed separately from other core analytics strategies. A truly integrated, centralized approach to AI and core analytics is the best place to start – and it requires tight alignment between hardware and software.

SAS and Intel have prepared for this moment for years, collaborating at the chip level to ensure that SAS® Viya® is finely tuned to take advantage of Intel capabilities.

Thanks to that collaboration, Intel chips are prepared for the processing challenges that come with delivering AI at scale, on premise or in the cloud. Leaders at organizations using **SAS capabilities** to deliver on their AI and analytics strategies don't have to wonder whether or not they're maximizing Intel capabilities. If they're using SAS, they are. It just works.

Here are 4 ways SAS® users benefit from deep SAS and Intel collaboration.

1

Advanced workload management

In an AI environment and at any given time, analytics solutions are focused on a wide range of tasks that stretch well beyond deep learning training – conducting data prep functions, running IO, and many other operations. Ultimately, this can force tradeoffs between performance and utilization.

Why not use a separate analytics appliance to focus solely on AI? It's tempting in theory, but in practice, this is an unwieldy, inefficient approach. Leaders in AI have found that it's far more practical to rely on a single, central, integrated analytics solution built to manage a wide range of AI and non-AI workloads. This also helps ensure that analytics capabilities are accessible to a wider range of users.

Some organizations have extremely high demands for specific workloads that are required time after time. In these cases, highly specialized, customized analytics tools may be warranted. Given the expense and ongoing resources required to support these types of tools and accelerators, they should be the exception rather than the rule.

SAS, running on Intel, has a role to play in each of these scenarios. It has both the extensive functionality and massively parallel processing to run the full range of AI and non-AI workloads. And in cases where specialized analytics capabilities are needed, it can be customized. The result is a more efficient approach to AI that opens the door to a wider range of tasks. With SAS and Intel, you can experiment more, remove the data science bottlenecks, and develop more accurate models faster, using familiar workflows. All with a lower total cost of ownership.

2

Tight collaboration on data management

Successful AI is preceded by a host of critically important – and familiar – data prep work. How will the data be managed? Transmitted? Stored? Organized? Cleansed? Addressing each of these activities beforehand means that data scientists and others can focus on the most important, productive, and valuable aspects of AI. But both research and anecdotal insights tell us that this more mundane prep work is still frequently ignored, resulting in data scientists spending the majority of their time focused on storing and cleaning data, rather than performing and managing AI tasks.

That's a problem. But the right solutions can help. Together, SAS and Intel have already built in the data management infrastructure and functionality needed to support AI processes, so users can focus far less of their time on these aspects of AI. This isn't just about compute capabilities, either. Networking, memory, storage, and storage acceleration (from Intel® Optane™) are all key parts of the infrastructure that come built in thanks to the SAS and Intel approach. As a result, users can spend more time focusing on the higher-reward, higher-value aspects of AI, confident that the data prep work is being executed in accordance with the highest standards.

Leaders in AI have found that it's far more practical to rely on a single, central, integrated analytics solution built to manage a wide range of AI and non-AI workloads. This also helps ensure that analytics capabilities are accessible to a wider range of users.

Here are 4 ways SAS® users benefit from deep SAS and Intel collaboration.

3

Diverse hardware and software capabilities for equally diverse AI challenges

AI challenges draw from many different data sources operating in a wide range of environments. When these hardware and software assets – edge devices, laptops, mobile devices, servers, workstations, you name it – are already optimized with AI capabilities, users can start seeing the impact of AI faster. IT leaders can also scale up these assets more quickly when the time is right. This gives you the ability to move code from modeling to production in a tight, streamlined manner. That's why Intel has focused on bringing AI capabilities to the full range of its products, all of which arrive with those capabilities already built in and ready to integrate.

Scale and fit are important parts of the equation. It's one thing to be able to check the box on AI capabilities on a device, but are they the right capabilities for the job the device is expected to perform? Are the AI capabilities outsized for the job at hand, or are they not enough? Are these capabilities delivered in hardware and other devices that are rugged enough to thrive even in challenging conditions – on the factory floor or in trains, truck fleets, or far-flung server rooms? Obviously, natural language processing for a stationary digital assistant device requires vastly different capabilities than an IoT-enabled AI solution for tracking the performance of a diesel-powered train engine. Not only does the analytics solution need to reflect this difference, but so do the data processing capabilities.

4

SAS® automatically taps into chip-level AI features

IT departments and business users executing AI strategies are likely relying on several different solutions – not just one. Using SAS as a platform, users can integrate these solutions in a centralized environment that allows them to deploy in a more disciplined, strategic manner. And since SAS is tightly integrated with Intel, they are equipped to extend the benefits of updates across the entire ecosystem of tools and solutions at the moment of availability. Intel is optimized to work with solutions like these, such as PyTorch and TensorFlow, as well as SAS and open source tools, so when they're all deployed in the same environment, the benefits are easy to tap into.

For example, several recent generations of CPUs have added instructions to accelerate computationally intensive work in analytics and AI, such as AVX (advanced vector extensions). This is an important feature for AI tasks that rely on matrixed and vector operations, and it can all be coordinated through SAS solutions. Future improvements in AMX (advanced matrix extensions), which allow users to process large amounts of data through matrix multiplications, will also be optimized for use in SAS and a number of complementary solutions, opening the door to dramatic performance improvements.

IT departments and business users executing AI strategies are likely relying on several different solutions – not just one. Using SAS as a platform, users can integrate these solutions in a centralized environment that allows them to deploy in a more disciplined, strategic manner.

Better together

Are AI and analytics separate endeavors? That's the real question at the heart of AI technology strategies. Technically, AI can be handled separately, with separate workflows and supporting technologies.

But should it be? For IT leaders actively grappling with the practical considerations of implementing new AI capabilities, the answer is to inject AI capabilities into existing environments – starting with analytics. This approach requires less effort (and investment) than standing up a separate analytics appliance dedicated to AI. And when it's done with SAS solutions, it takes advantage of critical chip-level optimizations from Intel and SAS collaboration. It's a smarter, more streamlined, performance-aware approach to AI – and it's ready today.

Right now, we're helping IT and business leaders across industries launch and operationalize their AI strategies, and we can help you. If you have questions about how to do it, let's talk.

To learn more, visit sas.com/intel

POWER OF THE PARTNER

At SAS, we love bold questions. And when we combine our analytics leadership with the innovative technology and expertise of our partners, we help our customers turn data into answers. That's the kind of curiosity that moves the world forward. That's the **Power of the Partner.**