

CACHING IN FOR EFFICIENCY

A CLOSER LOOK AT THE LATEST IN A CACHING TECHNOLOGY THAT DOUBLES APPLICATION PERFORMANCE WITH BEHAVIOR ANALYTICS

Is all caching technology created equal? This paper explores the vast differences in caching approaches and how the emergence of new engines is doubling application performance.

Understanding Caching

Latency can wreak havoc on applications running on your physical or virtual servers. Left unchecked, latency causes significant user frustration. This is one reason for all the focus in recent years on adding caching capabilities for performance while also reducing the amount of I/O traffic pushed to the backend storage subsystem.

With caching, the system stores data blocks in a way that ensures it can rapidly fulfill subsequent requests without having to constantly go back and forth between the server and storage. “It’s about making the best use of the resources available and providing the user with an optimal experience every time,” says Robert Woolery, senior vice president of product management and marketing with Burbank, California-based ConduSiv Technologies.

Conquering Cache Confusion

There are significant differences in caching, however, with a wide array of software providers and storage companies offering competing approaches. And, while the end goal is to eliminate latency and reduce I/O traffic, each approach impacts the operating environment differently, some leaving residual effects.

Understanding the difference between caching approaches starts with understanding exactly where the caching solution resides. For example, within the increasingly popular virtualized environment, there are the many layers of software that combine to form a hypervisor.

By design, the hypervisor allows for a collection of virtual machines in a layered fashion, with each behaving like a physical server. Today, most caching solutions either install beneath or within the hypervisor, or in some instances within the storage system.

IntelliMemory[™], the caching engine within ConduSiv’s V-locity[®] I/O optimization software, is a different story—it operates within the OS at the top of the technology stack, which gives it the advantage of application-aware caching. “This is an important distinction because we are located right where the I/O is created, meaning we have an immediate and ongoing opportunity to optimize based on application behavior,” says Woolery. “All other solutions reside farther down in the stack, after I/O streams from VMs have already been mixed and randomized. Without application awareness, they are forced to resort to predictive analysis to optimize caching performance, which often results in inaccurate guesses and steals unnecessary bandwidth. They make assumptions based on individual actions that the system will need I/O operations for blocks that are not

needed. This rapidly compounds into a serious issue and results in less than expected performance. There is only one way to take the ‘guess work’ out of caching, and that is with behavior analysis made possible via application awareness on the OS layer.”

V-locity leverages a built-in behavior analytics engine to accurately monitor hot data to see what’s common, determines use frequency, understands common access times, then caches server-side to minimize unnecessary I/O. V-locity can therefore cache within memory and hold blocks because it truly understands what the system needs. “V-locity sees requests from the application and, by the time the next request comes in, it already knows what is needed without wasting resources,” he says. “Relying on guesses is a problem that results in enormous amounts of I/O creation and wasted resources.”

According to Woolery, while caching is built-in to the Windows operating system as well as an assortment of business applications, it is inefficient. The caching technology in V-locity consistently proves to be at least 50% faster. “We understand how to make I/O more productive and faster based on understanding application behavior. Behavior analytics ensures the most effective caching method possible—reducing I/O traffic overall, which increases IOPS performance, reduces latency, and lowers overhead on shared storage by reducing physical read I/O requests on storage devices.” Woolery says, “It’s not just how fast an application interacts with the operating system, it’s about efficiently stripping away the latency associated with moving data to storage and back.”

V-locity is also able to avoid using a lot of CPU power. Because it works within the OS, it understands when applications make requests, when resources are available, and throttles back when necessary. “We never create a resource contention problem. We back off before a resource is needed,” he says. “We see a request and make sure we release all the memory we can to make sure the application has what it needs.”

V-locity is unique in that it is not an agent; instead, it’s a very thin filter driver. “The reason we can stay small is that we are essentially a consultant telling the operating system what it should do,” Woolery says. “We make the operating system more intelligent with recommendations based off application activity. It’s the OS that does all the work. We give hints on how to write and read data so the OS can improve its efficiency through improved I/O operations. How much improvement does this provide? With over 1,500 V-locity customers to back this claim - our typical customer sees 50% or more performance improvement.”

“IT IS NOT JUST HOW FAST AN APPLICATION INTERACTS WITH THE OPERATING SYSTEM, IT IS ABOUT EFFICIENTLY STRIPPING AWAY THE LATENCY ASSOCIATED WITH MOVING DATA TO STORAGE AND BACK.”

— ROBERT WOOLERY, Senior VP of Product Management and Marketing, ConduSiv Technologies.

Bottom Line

Simply put, it’s important to pay attention to how latency is being addressed through effective caching. V-locity is an I/O optimization solution designed to help protect infrastructure investments by improving productivity by 50 percent or more through a 100% software approach.

“This is a software-only approach to better performance. You can install and start receiving benefits within an hour,” Woolery says. “All without adding additional hardware.”

Click here to learn more about V-locity I/O optimization software, the free evaluation offer, and the built-in benefit analyzer to track results.