


CASE STUDY

Intel® Cache Acceleration Software

Performance: Data-Intensive Computing
Data Center Software Division



Accelerate Business Analytics Performance Cost Effectively with Intel® Cache Acceleration Software and Intel® Solid-State Drives (SSDs)

Intel® Cache Acceleration Software (Intel® CAS) with Intel® Solid-State Drives increases application performance with intelligent caching to enable up-to-date business intelligence while decreasing time and cost.

SAS, the global leader in high-performance analytics, enables customers to measure what matters most today, reveal best actions, expose threats, and gain predictive insights that compel the right actions. In a world where “big data” grows exponentially every day, SAS provides a suite of intuitive interfaces and the ability to deploy real-time analytics directly to mobile devices, so every decision maker—everywhere—can monitor key metrics and make informed decisions.

One challenge many companies face when monitoring this glut of structured and unstructured data is analyzing it and getting results in a timely fashion. Cheap, abundant hard disk drive (HDD) storage has made it easy to save all kinds of data, yet these HDDs are also a bottleneck when users try to retrieve the data for analysis.

All organizations with access to large data collections should try to harness the most relevant data and use it for better and faster decision-making.

Challenges

- **Reduced productivity.** Existing hard disk drives experience significant lag when querying data sets larger than five million records. This forces developers to spend additional time waiting for queries to load.
- **Cost of improving performance.** Lengthy query load times impact the quantity and timeliness of customer-ready data analysis. A full migration to SSD-based storage improves performance dramatically, but the hardware and time investment required may be prohibitive for some customers.

Solution

- **Intel® CAS and Intel® Solid-State Drives.** We suggest a seamless upgrade using Intel® CAS for SAS input and migrating the SAS WORK directory to an Intel® SSD. This approach improves performance cost-effectively—nearly twice as fast as HDD only...for about one-tenth the cost of a full SSD implementation.

Benefits

- **Nearly double performance with a minimum investment.** Intel® CAS paired with Intel® SSDs for both caching and SAS operations enables developers to perform queries in significantly less time, allowing them to complete their analyses faster and reduce costly load waiting time...all while using existing storage. This solution is significantly faster than HDD-only systems, yet costs about 10% of a total solid-state solution.



“SAS understands that customers need fast access to big data for making critical business decisions. Quickly getting the right information impacts the bottom line. Investment in the infrastructure that allows faster access to their data is necessary, but it doesn’t always have to break the bank. In joint testing with Intel, SAS discovered that Intel® Cache Acceleration Software improved data retrieval significantly at a reasonable cost. Intel® CAS will allow many of our cost-conscious customers to improve their analytics performance economically.”

— Cheryl Doninger, Senior Director
Research and Development, SAS



What is Intel® CAS and how does it work?

Intel® CAS is a server-side caching solution that supports directly attached SSDs and SSDs accessed through SAN storage protocols (such as iSCSI, Fibre Channel, FCoE, etc.), and provides a write-through cache for HDD-based storage solutions. Available on Windows* and Linux* platforms, **Intel® CAS integrates seamlessly, with no changes to your application or backend storage.**

As an application reads data from storage, Intel® CAS copies the data to the cache, which serves subsequent reads at higher speeds. Meanwhile, all data is written synchronously to backend storage and the cache.

Why Intel® CAS with Intel® Solid-State Drives?

Our High Endurance Technology (HET) provides 30 times the endurance of standard multilevel cell (MLC) technology. Intel® SSDs with HET provide considerable write endurance: Ten full drive writes per day over five years.

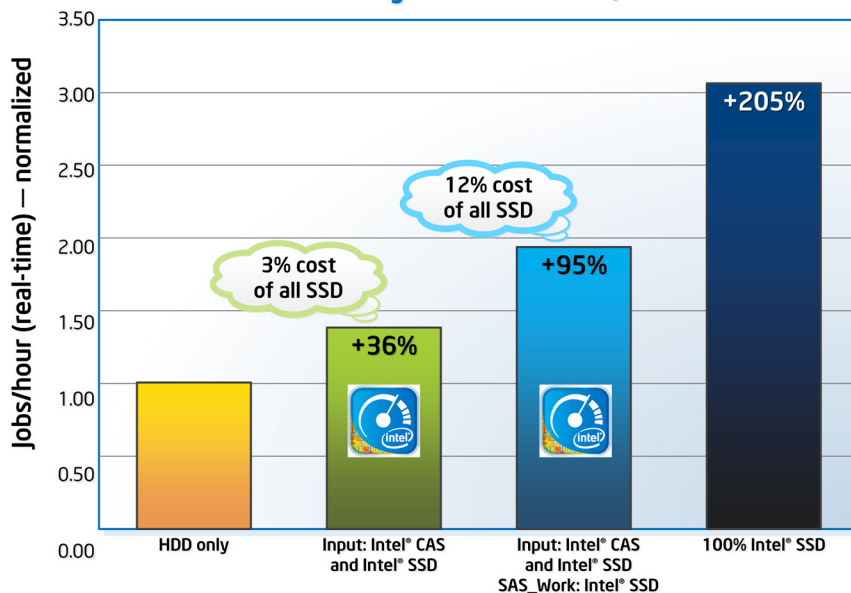
- **Intel® SSD 910 series.** Delivers extremely high performance storage with 180 K random reads and 75 K random writes.¹
- **Intel® SSD DC S3700 series.** Delivers data at a breakneck pace, with consistently low latencies² and tight IOPS distribution (75 K random read IOPS³).

Intel® CAS and Intel® SSD performance benefits

Intel and SAS jointly tested a typical 40-user SAS mixed analytics environment⁴ operating on an HDD-only system.⁵ With this as our baseline, we compared performance against the same system with three changes:

- Intel® CAS and Intel® SSD for caching input.⁶
- Intel® CAS and Intel® SSD for caching input, plus Intel® SSD for SAS WORK.⁷
- Intel® SSD for caching input, SAS work, and output.⁸

SAS mixed analytics workload, 40 users



As the graph above shows, adding Intel® CAS and an Intel® SSD for caching improves performance by 36%, yet costs less than a nickel on the dollar compared to a full SSD solution. Adding Intel® CAS plus Intel® SSDs at the cache and SAS operation layers **nearly doubles performance**, at one-tenth the cost of a full SSD solution.

1. Read and write IOPS (Full LBA Range, Iometer* Queue Depth 32 per NAND module). Default/max performance mode: Random 4 KB reads of up to 180,000 IOPS; random 4 KB writes of up to 75,000 IOPS.
2. 50 µs typ; <500 µs max: 100 GB 4 KB random writes QD=1 at 99.9% of the time across 100% span of the drive.
3. Random 4 KB reads.
4. **System configuration:** Generic "white box" w/ Sandy Bridge EP 4S (2.9 GHz) processor and 256 GB of RAM running RHEL 6.3. **Application:** Traditional SAS v9.3 mixed analytics workload in Linux OS environment.
5. HDD setup: Two LSI 620J enclosures, two 1078 MegaRAID SAS controllers, and twelve 500 GB Seagate Momentus HDDs.
6. Same as the HDD setup, with one 100 GB Intel® SSD DC S3700 and Intel® CAS for caching input.
7. Same as the HDD setup, with one 100 GB Intel® SSD DC S3700 and Intel® CAS for caching input, plus one 800 GB Intel® SSD DC S3700 for SAS_WORK.
8. Eleven 800 GB Intel® SSD DC S3700, with an onboard controller.

Spotlight on SAS

SAS is the leader in business analytics software and services, and the largest independent vendor in the business intelligence market. Through innovative solutions, SAS helps customers at more than 65,000 sites improve performance and deliver value by making better decisions faster. Since 1976, SAS has been giving customers around the world The Power to Know®.


SAS leads the way in high-performance analytics, helping customers tackle complex problems using big data, and gain highly precise insights to speed information and outperform competitors. SAS customers gain insights in minutes or seconds that once took days or weeks!

SAS is headquartered in Cary, NC. For more information, visit www.sas.com.

Learn how Intel® CAS seamlessly accelerates applications and more at intel.com/cas.

Find the Intel® Solid-State Drive that's right for you. Visit intel.com/go/ssd for product information.

To learn more about other Intel business solutions, please visit the Reference Room at intel.com/references.



Intel does not control or audit the design or implementation of third-party benchmark data or websites referenced in this document. Intel encourages all of its customers to visit the referenced websites or others where similar performance benchmark data is reported and confirm whether the referenced benchmark data is accurate and reflects performance of systems available for purchase.

Copyright © 2014 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others. 330289-002EN