



CASE STUDY

Powering AI Modeling for a \$24 Billion Investment Portfolio



Accelerated
time-to-market for
new trading ideas

10x to 20x

Processing improvement
for quantitative
researchers



Provides on-demand
path for future growth
with highly scalable
architecture

“Pure FlashBlade gives us significant headroom over what we need for the business today.”

JAMES BLACKBURN,
HEAD OF DATA ENGINEERING, MAN AHL

Man AHL is a pioneer in the field of systematic quantitative investing. With more than \$24 billion in assets under management, the firm uses mathematical models to make all its investment decisions. The firm adopted Pure Storage to deliver the massive throughput and scalability required for its most demanding simulation applications.

Performance and Productivity Optimized for Quants and Technologists

Given its business model, Man AHL has attracted two extremely valuable sets of resources—the roughly 50 quantitative researchers (“quants”) who formulate the ideas driving the investment models, and the 60-plus technologists who both develop the quantitative research platform and turn quant ideas into trading strategies that can be executed by computer. One of the firm’s highest priorities is delivering a robust and ever-more-performant technology platform as a means of optimizing the performance and productivity of its quants and technologists.

In 1987, stock markets around the globe suffered their greatest one-day losses ever. That also was the year the investment management firm AHL (now Man AHL) was founded to deliver a revolutionary approach to hedge fund management. The firm uses mathematical models (algorithms) to make all its investment decisions, which are executed solely by computer without human involvement.

“Our technology is key in giving the quants an environment in which they can explore and test their ideas, as well as a production environment that will faithfully implement the models they create,” says Gary Collier, CTO of Man Group Alpha Technology in London. Man AHL is part of Man Group, whose five investment management businesses manage over \$113 billion as of mid-2018.

“ We’re an extremely agile business and need to be able to add terabytes of storage almost on a whim.”

**GARY COLLIER,
CTO, MAN GROUP ALPHA TECHNOLOGY**

Developers, Analysts Require High-Performance Storage

The demands on technology infrastructure at the firm are significant. “In the development of our trading models, we need to look at large amounts of data to run experiments and sift through all manner of possible ideas that might make money for our investors. Only those models in which we have sufficient conviction will then be put to work in the production trading environment,” Collier says.

In particular, the use of machine learning techniques is an evolving area within quantitative investment strategies, and Man AHL has been actively applying these techniques within client programs since 2014. One of the broadest potential benefits of machine learning to Man AHL’s investors derives from its ability to handle large volumes of data. Graham Robertson, Head of Client Portfolio Management for the firm, recently wrote: “If this is the age of the data deluge, then machine learning algorithms have the potential to dramatically increase investors’ ability to process and analyze information on markets.”

In 2016, the technologists at Man AHL had concerns about the ability of its legacy storage infrastructure to meet the demands of the quants and developers. “We had an HPC (high-performance computing) environment with shared storage that supported all our research and software development activity,” noted James Blackburn, Head of Data Engineering at Man AHL.



“At times, you would have people manipulating large numbers of files at the same time quants were trying to run large simulations. Our storage system just didn’t have the IOPS to support all those demands simultaneously, which meant some people would get frustrated by the response times. We needed something that would provide much greater IOPS and bandwidth.” Collier noted, “We’ve had poor experiences in the past with expensive—and supposedly highly parallel—storage solutions simply not being able to deliver the bandwidth and IOPS we require. This time, we wanted to make sure we conducted extreme due diligence to explore every possibility.” Blackburn was charged with managing the search for a new storage infrastructure.

In addition to high IOPS and bandwidth, “We also wanted the new storage to be highly scalable,” Collier added. “We’re an extremely agile business and need to be able to add terabytes of storage almost on a whim. We just want to be able to start where we are, then add capacity and IOPS as necessary. In other words, we want something that scales linearly.”

In addition, Collier noted, “We do not want to dedicate multiple engineers to the job of building and maintaining bespoke storage solutions, or to administering complex-to-manage vendor products.”

Proving the Concept for Modern Analytics

Knowing that their need for high bandwidth and IOPS meant next-generation storage, the team investigated suppliers of all-flash products. “Pure Storage is very well-respected, both by organizations like Gartner and by users,” Collier says. “The customers who have purchased and deployed Pure take the time to write Internet posts about their positive experiences. That was a good sign for us from the start.”

Based on its initial due diligence, the Man AHL team arranged for a proof-of-concept trial of a Pure FlashBlade system from Pure Storage. FlashBlade is the industry’s first data hub to consolidate shared file and object storage on a single scale-out platform that leverages Purity software’s Variable Block Metadata Engine and intelligent load balancing to accelerate data-intensive applications. Given the enormity of the financial simulations conducted by Man AHL, FlashBlade’s massively parallel storage capabilities have proven to be a match for these demanding workloads.

Man AHL’s FlashBlade configuration has eight blades, each with 52TB raw capacity. Factoring in the overhead of the Purity OS and erasure coding, Man AHL is achieving 329.14 TB of effective capacity thanks to an aggregate data reduction of 1.5:1 for all workloads running on FlashBlade.

“In the proof-of-concept phase, we pushed the Pure Storage system to the max,” Blackburn reported. “We tested total throughput for large I/O, maximum IOPS, and other metrics. Overall, the system scaled as promised. We achieved 6GBs bandwidth for large I/O reads and 3GBs for writes. And the 600,000 IOPS we achieved were far more than we ever had before.”

“ FlashBlade gives us significant headroom beyond what we need for the business today. This gives us confidence knowing we can simply add a blade to gain more capacity and preserve the investment we already have.”

JAMES BLACKBURN,
HEAD OF DATA ENGINEERING, MAN AHL



Performance Improvements Create Time-To-Market Advantage

The impact of the new data hub was noticed immediately. “Many of our researchers found that the introduction of FlashBlade has made it easier to use Spark for performing multiple simulations. One of them experienced a 10x to 20x improvement in throughput for his Spark workloads compared to the previous storage system,” Blackburn says.

Added Collier: “Our quants want to test a model, get the results, and then test another one, and another one—all day long. So, a 10x to 20x improvement in performance can be a game-changer when it comes to creating a time-to-market advantage for us.”

The greatest benefit from installing Pure FlashBlade, Collier says, “is significantly improved productivity for the team and accelerated time-to-market for new trading ideas.”

The POC involved workloads from across the firm’s research teams, “and we picked jobs that were very demanding on I/O, like our Jenkins builds and Spark users,” Blackburn says. “As we progressed, we loaded more and more concurrent workloads and did benchmark testing to see how close we were to the limits of the system. What we found is that FlashBlade gives us significant headroom beyond what we need for the business today. This gives us confidence knowing we can simply add a blade to gain more capacity and preserve the investment we already have.”

A month into the POC, the firm’s entire research department had been migrated to FlashBlade. In addition, Man AHL runs a concurrent 50 TB Mongo database containing order-book data from stock exchanges and other sources. “During the POC, we put a replica set of the Mongo data on FlashBlade using an NFS mount. It worked so well that we left it there,” Collier says.

Beyond meeting the high-performance requirements of its core simulation applications, Blackburn says, “the real benefit of the FlashBlade is the consistent I/O and metadata access. To have a shared substrate of I/O that is consistent throughout the day makes users very happy. Storage used to be the bottleneck, and now that has been removed.”

“ If we extrapolate the way we have been growing, we double our storage needs about every 18 months. Having something as flexible and expandable as FlashBlade is what our business demands.”

GARY COLLIER,

CTO, MAN GROUP ALPHA TECHNOLOGY



GEO

EMEA

INDUSTRY

Banking and Financial Services

SOLUTION AREAActivate Real-time Analytics
Power Artificial Intelligence**PURE STORAGE®
PRODUCTS IN USE**

FlashBlade®

ABOUT CUSTOMER

In 1987, stock markets around the globe suffered their greatest one-day losses ever. That also was the year the investment management firm AHL (now Man AHL) was founded to deliver a revolutionary approach to hedge fund management.

www.ahl.com**Effortless Management and a Clear Path for The Future**

Man AHL benefits from more than FlashBlade's raw performance. "There are significant savings from the simplicity of operation. It just works," Collier noted. "The fact that Pure Storage is plug-and-play adds to our operational efficiency. We've essentially cut our headcount for storage management from one-and-a-half to the equivalent of a half-time role. And going forward it'll approach zero. That means we can redeploy those assets to higher-priority activities."

Looking to the future, Collier says FlashBlade is a perfect fit for the firm's growing—but at times unpredictable—storage and computing needs. "If we extrapolate the way we have been growing, we double our storage needs about every 18 months. And that's a scary prospect. FlashBlade is very compact and scales easily just by adding blades. Having something as flexible and expandable as FlashBlade is what our business demands."

Another attractive long-term benefit is Pure Storage support for the OpenStack initiative. "We are heavy users of Docker," Collier says, "so we plan to take full advantage of OpenStack and Docker integration on the Pure data platform. That makes it very easy from a DevOps standpoint to create Docker volumes and attach them to containers."

Above and beyond product satisfaction, the Man AHL team has benefitted from the extreme attentiveness of Pure Storage customer and technical support. As early as the POC, Blackburn says, "we saw how responsive the Pure team is. We had answers to questions generally within an hour of sending them, whereas we find other vendors can take weeks to respond. We've had astonishing turnaround from Pure."

Important Information

Opinions expressed are those of the author and may not be shared by all personnel of Man Group plc ('Man'). These opinions are subject to change without notice, are for information purposes only and do not constitute an offer or invitation to make an investment in any financial instrument or in any product to which any member of Man's group of companies provides investment advisory or any other services.

