INNOVATE WITH DATA IN INVESTMENTS

A DATA-CENTRIC APPROACH TO FASTER AND BETTER INVESTING

CONTENTS

PURESTORAGE®

UNPRECEDENTED CHALLENGES

DIGITAL TRANSFORMATION

Investment firms (investment banking, asset management and wealth management) are going through a digital transformation driven by a more dynamic investing environment, increased customer expectations, stricter regulation and fierce competition from FinTechs.

What does digital transformation look like in investing?

For the buy-side it may mean:

- being able to quickly integrate and analyze both market data and alternative data to get an investing "edge" and enable "real-time" investing decisions
- using emerging technologies such as Artificial Intelligence (AI), to create lower-fee, passive investing strategies as a response to competition from robo-advisors
- being able to combine machine learning algorithms with predictive analytics and relevant data, to help financial advisors generate timely, customized advice for clients. It's about fast and smart data driving higher investment returns.

For the sell-side, digital transformation may mean:

- being able to leverage digital platforms to provide services in high volumes, spreading out fixed costs
- using technology and data analysis to better understand their clients and find new ways to cross-sell and improve relationship stickiness (client-centric vs product-centric engagement model); improve client experience
- optimizing back-office processes to cut costs and achieve higher level of services (e.g. shortening settlement time T+2, client on-boarding)
- using AI/Machine Learning (ML) and a variety of data sources to create better risk models for running the business and for compliance (e.g. stress testing).

For all the players, digital transformation means using data, analytics and Artificial Intelligence to become more proactive in critical areas such as cybersecurity, fraud management, client interaction and overall experience.

BARRIERS TO DIGITAL TRANSFORMATION

A fundamental barrier to achieving digital transformation in investments is often the complexity that has accumulated in silos. For many large firms, their existing data strategy has become immersed in a series of tactically extended, high latency silos that inhibit data-sharing, or timely access to data.

The data sources are disparate and may reside in the front office (for example, customer services, quantitative research, sales, trading), middle office (including risk management, compliance, trade surveillance, underwriting, merger and acquisition due diligence) or back office (such as operations, settlements or accounting).

There are consequences in missed trading opportunities, frustrated quantitative research, or a declining customer experience leading to loss of trust. This is happening at a time when technology-driven start-ups are disrupting the business model using leading-edge analytics, AI/ML/Deep Learning (DL) and a highly automated operational approach.

DATA-CENTRIC TRANSFORMATION

Consolidating, connecting and accelerating data can overcome many of these challenges. Designing for the greater value of data correlation across use cases results in the elimination of legacy silos and facilitates the flow of data into analytics for near-real time or real-time insights. Connecting multiple, diverse data sources and being able to correlate and analyze the data at scale leads to more frequent, more current and deeper insights.

Having a data-centric platform with scalable processing power and capacity opens the door to building new, AI-based services, which lead to better customer facing services, faster risk assessment, better cybersecurity and faster fraud detection and prevention.

USE CASES AND SOLUTIONS

We will examine a number of use cases where Pure Storage® can help with establishing a more agile, competitive investment operation, and provide the business case for a modern, data-centric architectural design.

ACCELERATING QUANTITATIVE RESEARCH

Quant Challenges

Quantitative (quant) research includes development, testing and deployment of models that search for new investment strategies. To establish and test an innovative investment hypothesis, quants use historic time-series data enriched with "alternative data." Alternative data can take many forms such as logistics data, foot traffic, sentiment data, satellite images, geo location, demographics, social media, which is unstructured and generally high volume. Legacy storage systems do not offer the performance needed at this scale – the larger the data set, the longer it would take to process.

Pure Storage for Quants

A data hub built on Pure Storage enables quants to process more data significantly faster, enabling quicker time to value for the firm and the client. They can explore larger data sets, collaborate, and back-test models far more rapidly than conventional systems.

> *"Our quants want to test a model, get the results, and then test another one, and another one – all day long. So, a **10x-20x improvement in performance** can be a game-changer when it comes to creating a time-to-market advantage for us."*
>
> **– Co-CTO, MAN AHL**

As data sets become larger, analysts need assistance from ML algorithms to sift through the volumes of material. The Pure Storage all-flash arrays can sustain performance for multiple workloads, such as massively parallel tools like Kafka and Spark. At the same time, operationally, if more capacity or processing performance is required, then it can be added independently, with no downtime, in a true "scale-out" fashion.

TRANSFORMING THE CLIENT EXPERIENCE

Sell-Side Challenges

In the post-crisis era, the sell-side has struggled with lower fees and disintermediation. At the same time, the buy-side has developed technological muscle, and is demanding more customized services from the sell-side. To create these differentiated services, investment banks can tap into their huge volumes of transaction data. The challenge is that data is kept in silos by lines of business, or segregated into hot-warm-cold tiers. This happens largely because it was cost-prohibitive to keep all data available, all the time on legacy platforms. The limitation can negatively impact the effective use of advanced analytics and the speed and quality of new services development and delivery.

Pure Storage for the Sell-Side

Keeping all investment data on a Pure data storage platform reduces latency for many workloads: from pre-trade and client facing portfolio analytics to post-trade processing. In addition, Pure Storage all-flash arrays provide greater consistency, and accelerate the production of daily client investment reports, which significantly improves client relationships, and maintains consistent compliance with client Service Level Agreements (SLAs). All-flash arrays can also provide the performance needed by AI powered applications.

GOVERNANCE, RISK AND COMPLIANCE (GRC)

Compliance Challenges

The complexity and cost of regulatory compliance and the associated risk management is growing worldwide. Investment firms are constantly reviewing how to reduce the cost and burden of compliance, while minimizing risk. A key compliance challenge is to keep up with the ever more stringent requirements of regulations such as CCAR (Comprehensive Capital Analysis and Review), MiFID II (Markets in Financial Instruments Directive) and Basel BCBS239 reports, EU General Data Protection Regulation (GDPR), plus Anti-Money Laundering (AML) and Know Your Customer (KYC) suspicious activity detection and prevention.

The reporting process can be complex and demanding faster access to large data sets. For example, MiFIDII requires the sell-side to undertake new, real-time reporting obligations, including the need for capturing and recording all methods of communication surrounding price negotiation. For many firms this sales and trader workflow requirement will stress traditional data storage systems, particularly during heavy trading periods.

CCAR in the US is an annual exercise which requires a more accurate assessment of risk and includes a stress-test process to make sure that the bank has sufficient capital to survive a crisis. Consequently, it is vital that the investment bank data platform has both the capacity and processing power to run multiple iterations on these large data sets.

GDPR, although an EU regulation, impacts financial firms globally and it requires strict personal data governance: protection, privacy, portability and the right to erase. More than half (54%) of firms surveyed by IDC admitted encryption and pseudonymization of data as GDPR requirements pose the "greatest challenge" to them.

Pure Storage for GRC (including Cybersecurity)

By 2021, 45% of all banks will have invested in automated GRC applications to improve operational performance, and substantially reduce the operating expense associated with manual processes**. In support of this transformation, data from different sources, and in a variety of formats, needs to be integrated and shared so it can be accessed quickly by analytics and AI workloads.

The volume of data that is classified as sensitive and may require security consideration has been increasing due to new and updated privacy regulations like PCI-DSS, GDPR. This challenge is only going to get worse, since regulations are constantly changing and privacy concerns are escalating. It is advisable to be ready to treat more and more data as sensitive.

Pure Storage consolidates, connects and accelerates data for GRC applications, while keeping it secure at scale. With FIPS 140-2 validated AES-256 data-at-rest encryption (the U.S. government's highest security standard) and Rapid Data Lock technology, Pure Storage makes sure that sensitive data remains private and secure. Pure Storage complements this ultra-high level of security with automated encryption key management.

FlashArray //M

To top it off, Pure has integrated with Vormetric to offer data compression on encrypted data, which is a first in the industry. Users no longer need to make the choice between data security or storage efficiency. As the volume of data grows exponentially and regulations are being updated to protect the data (e.g. GDPR which took effect in 2018), financial firms have been paying a high price (economically or in performance) to encrypt the data: encrypted data can consume 5x more storage than unencrypted data that has been reduced. Now they have a solution in the Vormetric Transparent Encryption with Pure Storage FlashArray™, which, in a recent test, has achieved a 4.8:1 data reduction ratio.

Pure Storage FlashBlade™ is an ideal hub for aggregating data from different sources and in a variety of formats, to streamline cybersecurity and compliance. Many Pure customers are using FlashBlade as the repository for logs from applications and devices. Log analytics tools such as Splunk, Elastic Search, Security Onion, benefit from fast access to all logs in one place. The log data tends to grow exponentially and FlashBlade offers a cost-efficient way to scale, even compared with cloud. Splunk customers have now the option to use SmartStore, which uses FlashBlade as an object store for logs.

For investment firms considering AI/ML/DL for GRC, the Pure Storage FlashStack for AI is a good way to jump-start the project. Being a pre-validated, pre-integrated solution optimized for AI workloads, it offers predictable performance and shortens the testing and implementation cycle. It is built on Pure Storage FlashBlade and Cisco's Unified Compute System (UCS).

***IDC FutureScape: Worldwide Financial Services 2019 Predictions, December, 2018*

ACCELERATING DEV/OPS

Dev/Ops Challenges

An investment firm is reliant upon dev/ops for its digital transformation. Overall, the aim is to reduce time-to-market for new application development and deployment by enabling parallel operations, collaboration and testing, either on-premises or in the cloud.

Investment firms are moving away from monolithic applications to modular apps and microservices. These new apps are typically built upon containers, which decouple the application from the infrastructure to make them more portable. The challenge is that containerized applications want to consume "storage-as-a-service", but legacy storage systems do not operate in this manner. Containerized environments are also highly fluid and need rapid scalability to 1000s of containers, which is again unrealistic for a legacy system.

Pure Storage For Dev/Ops

The Pure Storage all-flash arrays support multiple concurrent dev/ops projects with predictable performance. Pure Storage rapid restore capability allows developers to create multiple instances of a database in a fraction of the time it would take with legacy systems.

> *"With **Pure Storage**, we can **deliver new development environments in minutes, not hours or days**."*
>
> **– Vice President of Production Operations, eMoney Advisor (Part of Fidelity Investments)**

The Pure Storage Service Orchestrator (PSO) is designed to provide developers with the same experience gained from the public cloud. PSO integrates seamlessly with the container orchestration frameworks (Kubernetes), and functions as the control plane virtualization layer, which enables a containerized environment to move from consuming "storage-as-a-device" to consuming "storage-as-a-service".

OPERATIONAL EXCELLENCE

Operational Challenges

Legacy storage systems have complex ongoing operational needs which translate into a higher Total Cost of Ownership (TCO). Maintenance costs escalate towards the end of a lifecycle, and force the investment bank into "fork-lift" upgrades that are costly and interrupt business as usual.

Data storage infrastructure simplification is the foundation for increased IT agility and one of the enablers of growth in an investment banking business. In addition to simplification, there is infrastructure cost reduction as an ever present pressure, especially with data center space at a premium.

Pure Storage for IT Operations

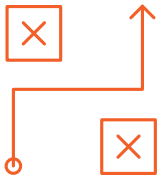

From its inception, Pure Storage had the goal of simplifying storage management, to reduce TCO and increase availability. Today, its software, Pure1®, includes an AI engine, Pure1 Meta®, which uses machine learning to make workload predictions. Pure1 continuously monitors rich telemetric data from 1000s of cloud-connected arrays – now more than 1 trillion data points a day, and more than 7 PB of total data. It's the analysis of this data by Pure1 Meta™ that enables Pure1® Support to predict vulnerability to known issues and proactively alert customers before problems develop.

To solve the storage life-cycle management challenge, Pure has introduced the Evergreen™ Storage Program, which helps customers future-proof their storage investment in Pure technology.

It is a subscription-based program with 2 levels: Gold and Silver. Evergreen Gold is the subscription most customers purchase, which includes the program components such as Free Every Three and Upgrade Flex for controller modernization as well as Capacity Consolidation for media modernization.

*"**Pure Storage Evergreen™** provides new controllers every three years, ensuring the hardware is modern and up to date. **This happens seamlessly, and while online, with zero down time.**"*

– Director of Technology, KordaMentha

Evergreen™ Storage

Pure Storage's Evergreen Storage Program

TECHNOLOGY ACQUISITION	ONGOING MANAGEMENT	TECHNOLOGY REFRESH*
Love your storage	Flat and fair	Free *every* three*
All-Inclusive software	Evergreen maintenance	Upgrade flex*
Right-size guarantee	White glove support	Capacity consolidation

LIFE-CYCLE INVESTMENT PROTECTION*

* Requires Evergreen Gold Subscription

See what IDC has to say about the Evergreen program, in this whitepaper.

The Evergreen Storage Service (ES2) is a private cloud storage service designed for no balance sheet impact under new accounting guidelines, enabling OPEX treatment regardless of term length*. With ES2, customers only pay for used effective capacity (measured daily), not provisioned capacity. The result is up to 2x greater cost efficiency compared to provisioned approaches.

Pure Storage for the Cloud

Many investment firms have a "cloud-first" initiative. For new applications this makes a lot of sense, especially if they are cloud-native. For software development the agility of the public cloud makes sense, but what if the application then needs to be moved to the private cloud because it needs access to sensitive or regulated data?

The reality is that a lot of critical back-office applications are still residing in silos, on-premises, in co-location or in hosted virtual private clouds. The future of cloud is hybrid cloud/multi-cloud, where applications, services, data and users can reside anywhere and stay connected.

Pure Storage's hybrid cloud solutions enable data sharing, application mobility, seamless management and automation across on-premises, hosted and public cloud resources. Cloud Block Store – industrial-strength cloud block storage, delivered natively in the cloud – lets you run existing enterprise apps in the cloud without rearchitecting anything. Develop your apps once – and run them anywhere.

The time that it takes to back up or restore data can become a business challenge. This challenge is amplified in the cloud world, where storage and bandwidth costs compound with data transfer rates. ObjectEngine taps the power of flash and cloud to deliver unprecedented restore performance on-premises while extending to public cloud for 11x9s of durability and pay-as-you-go economics. Best of all, it just works with your Veritas, Veeam, or Commvault workflow

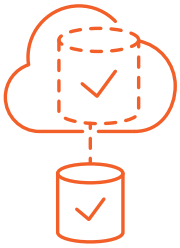

**OPEX treatment is subject to customer's auditor review.*

BUSINESS CASE FOR ALL-FLASH STORAGE INVESTMENT

The business case for all-flash storage investments centers on facilitating and accelerating digital transformation across all aspects of the investment business, from the front-office to the back-office. The metrics are based upon actual Pure Storage case studies.

Business Issue	Business Benefits from Pure Storage All-Flash	Metrics
Accelerating Quantitative Research	Quants can iterate faster the dev/test cycles for models and can use larger data sets, to combine historical market data and alternative data.	A 10x to 20x increase in quant data processing performance including data enrichment.
Transforming the client experience	Faster applications and faster access to data means a better experience for clients. Lower risk of breaking client SLAs.	<1ms latency. 4x more portfolio managers. 3x improvement in the client SLA operational deadline margin with greater consistency.
Efficient and Agile GRC (including Cybersecurity)	Encryption is built-in, facilitating data privacy and governance. Breaking data silos streamlines compliance reporting. High volume of data can be stored efficiently online for the long term and can still be available for analysis and reporting. Implement automation and AI faster.	Performance stays the same even with encryption enabled. Security certifications: FIPS 140-2, Common Criteria. Ability to compress encrypted data.
Accelerating Innovation	Reduced time-to-market for new application development and deployment leads to faster innovation for gaining competitive advantage. Ability to implement AI technologies faster.	Developers can spin up a database instance in minutes vs. hours or days with legacy systems. Data scientists can execute 4x to 5x the number of projects than previously using proprietary or self-build solutions.
Operational Excellence	Efficient use of data center space and power. Lower TCO for storage. Guaranteed access to the latest all-flash storage technology upgrades without any interruption to business-as-usual. Efficient and cost-effective scaling up in case of acquisitions. Choice of Storage-as-a-service ownership model, no balance sheet impact	Data center space was reduced by 79%. Power consumption reduced by 57%. Data-reduction rates of 5:1. Savings of $1.7M in just 6 months. ES2: 2x greater cost efficiency compared to provisioned approaches.

BUSINESS CASE STUDIES

Investment banks need to grow business and generate higher margins, in some cases motivated by the threat of new trading start-ups with significant fintech infrastructures or alliances. This selection of Pure Storage case studies emphasizes the real benefits being achieved for traditional investment banks on their way to digital transformation.

ACCELERATING QUANT RESEARCH

MAN AHL is a London based, diversified quantitative investment manager, with more than $19 billion in assets under management as of June 2017. “Quantitative” means all of the MAN AHL investment decisions are based on mathematical models or algorithms, using decisions made solely by computer, without human involvement.

10x-20x IMPROVEMENT
IN PROCESSING

The firm has adopted a Pure Storage FlashBlade solution to deliver the volume of stored data throughput and scalability required to meet the most demanding simulation applications. The highly scalable FlashBlade architecture provides an on-demand path to enable MAN AHL quantitative researchers (quants) to benefit from a 10x-20x processing improvement, and establish a foundation for future capacity growth.

Liquidnet

REAL-TIME ANALYTICS FOR TRADERS

Liquidnet connects nearly 1,000 institutional investors to large-scale trading opportunities in 46 equity markets across the globe. Nowhere is the adage “time is money” more true than in the world of institutional investing, where money can be made or lost on the basis of split-second decisions.

“A lot of the data that we collect has historically been used for post-trade analytics, but as the investment environment has become more complex and faster-paced, ***we want to provide traders more real-time analytics that will allow them to make pre-trade decisions****.”*

– Global Head of Product Support, Liquidnet

2.34 MILLION
EVENTS PER SECOND

The IT team saw a good match between the capabilities of FlashBlade and Liquidnet’s evolving needs, which center around making analytics available to traders in as close to real time as possible. During pre-production evaluation, FlashBlade was tested for its ability to handle Kafka stream-processing workloads, and the results startled the application team: 2.34 million events per second, compared to a significantly lower number of events for the legacy environment.

“FlashBlade had capabilities we thought were vital: ***the ability to process large volumes of streaming data in real time, and the ability to make a data set available to multiple workloads simultaneously.****”*

– Global Head of Product Support, Liquidnet

eMoney

TRANSFORMING THE CLIENT EXPERIENCE

eMoney Advisor is a cloud-based software-as-a-service (SaaS) platform that combines financial planning software, data aggregation, business analytics and other functions to help financial professionals build stronger client relationships, streamline business operations and drive revenue growth. They were acquired by Fidelity Investments in 2015. eMoney started with using Pure Storage FlashArray to accelerate their back-office databases. They were impressed with the performance but also with the reliability and availability of the systems. They went through three generations of FlashArray without disrupting operations.

100% AVAILABILITY
AVAILABILITY

"The stability, reliability and performance of Pure Storage arrays means we can ***raise our standards and deliver on our promises – even to the point of delivering 100% availability.****"*

– Vice President of Production Operations, eMoney Advisor (Part of Fidelity Investments)

Then they noticed that their database restore times were going up, due to increased volume of data. To be able to perform restores within their existing SLAs, they brought in Pure Storage's latest array, FlashBlade. In combination with Pure snapshots and Cohesity software, this solution provided multiple options for rapid restore/recovery.

Rapid restore is also used by software developers to quickly spin up copies of production databases. This capability reduced the dev/ops cycle significantly, contributing to faster time to market for new features, which, in turn, contributed to enhancing the customer experience.

"With Pure's snapshotting capabilities and high rates of data reduction, ***we can deliver new environments in minutes, not hours or days. Our DBAs can refresh databases quickly and easily, just with a click.****"*

– Vice President of Production Operations, eMoney Advisor (Part of Fidelity Investments)

FIRST 1R RATE

ACCELERATING CLIENT REPORTING

First Rate delivers investment performance reporting at the start of each business day to clients. When unexplained delays started occurring in the preparation of reports for one client, First Rate needed a solution. The answer was Pure Storage All-Flash Arrays, which has allowed First Rate to consistently meet the standards needed to provide timely strategic reporting, while also slashing operating costs.

Thanks to greater consistency of on-time availability for critical investment performance reports, client satisfaction has increased. In addition, SLAs have been strengthened, and are now offered to more clients thanks to vastly improved and predictable storage performance.

5:1 REDUCTION
IN DATA

The high level of IOPS, and consistently sub-millisecond latency, means overnight processing jobs are finished reliably and consistently, within SLA limits. In addition, the overall data-reduction rate of 5:1 results in lower data-center co-location costs.

Investec Asset Management | Out of the Ordinary®

ACCELERATING CRITICAL DATABASE

Investec is an investment management company based in South Africa. Because the trading window for business between South Africa and London is only two hours, Investec must swiftly execute batch jobs within a narrow time band to maximize market performance.

To reliably manage timely batch execution, Investec determined that all-flash storage was needed to accelerate the SQL database, which would directly improve performance and profitability. As a consequence of applying Pure Storage all-flash array, they were able to reduce all database job latency to below 1ms, and simultaneously increase by 4x the number of portfolio managers supported on the application.

80% REDUCTION
IN HARDWARE

From an infrastructure perspective there were additional gains. The Pure Storage all-flash array investment represented an 80% reduction in hardware, which resulted in additional cost savings for maintenance, cooling and power. Physical servers were no longer required for application data, and the storage hardware space reduced from 1.5 racks to only 10U.

There was a similar impact when the investment firm deployed a second Pure Storage all-flash array at a disaster recovery site, where the space demand reduced from 75% of a cabinet to only 8U.

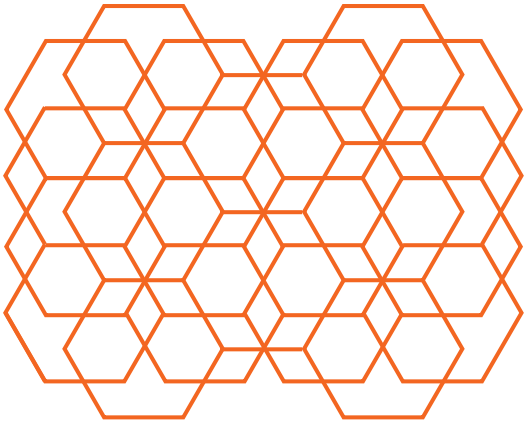

NEXT STEPS

Pure Storage offers expert guidance to assist investment firms in understanding, and then deciding if an **all-flash array strategy** is the appropriate storage solution for the business challenges of implementing a data-centric approach and leveraging AI and ML in the business model for greater productivity.

DATA-CENTRIC STRATEGY: SUBJECT MATTER EXPERT BRIEFING

An informal investments stakeholder briefing that is facilitated by a Pure Storage subject matter expert. The aim is to clarify the challenges, and potential solutions, for applying a best practice, data-centric storage strategy to protect and grow existing customer business, plus accelerate new client on-boarding. Please contact Pure Storage to arrange the briefing.

Contact	Pure Storage Inc., 650 Castro St #400, Mountain View, CA 94041, USA.
Visit	**www.purestorage.com**
Email	info@purestorage.com
Telephone	1 833 371 7873
Facebook	facebook.com/PureStorage
Twitter	twitter.com/PureStorage

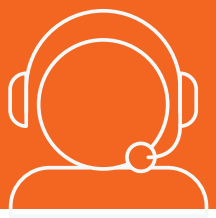

PURESTORAGE®