



Survey Highlights Hedge Firm IT Challenges in an Era of Constant and Rapid Change

IT managers at hedge funds are under pressure to meet demands for speed and processing power that lead to actionable insights while also keeping data safe and costs under control.

We are living in an age of surprises. As world events take unexpected turns they present opportunities for hedge funds that can act quickly enough to take advantage of them. And funds that incorporate the increasingly broad types and higher volumes of data into their pipelines—not just financial data but also factors such as real estate prices, weather patterns and social media—can develop higher resolution trade formulas to gain an even greater competitive advantage.

IT managers at hedge funds are under increasing pressure to meet demands for speed and processing power that lead to actionable insights while also keeping data safe and costs under control. It's a tricky balance. Should all processing be kept on premises? Rent processing or data storage in the cloud? A bit of both?

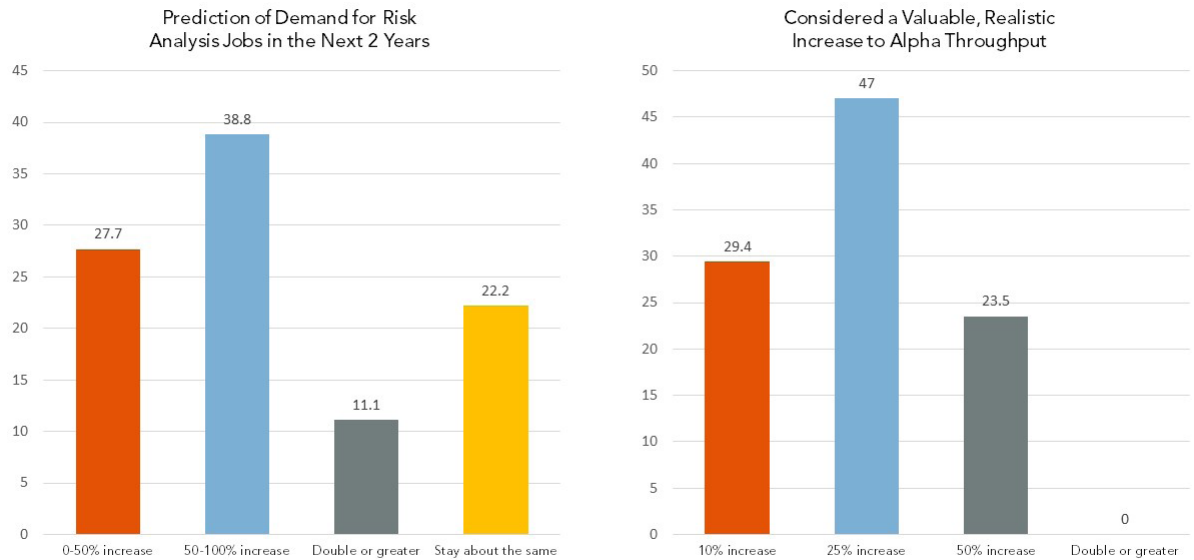
To gain insight into how hedge fund IT departments are dealing with these challenges, Avere Systems conducted a quick poll of the industry's IT professionals last fall. The results give a sense of financial services IT professionals' day-to-day concerns and needs.¹

The need to process more data, faster

The results indicate that IT managers are struggling to meet demand for more and more analytic processing power. A majority of respondents (78 percent) expect demand for risk-analysis jobs to increase in the next two years. The largest subset—39 percent—predict demand to rise by 50 to 100 percent. Twenty-eight percent think it will increase by up to 49 percent, and 11 percent think it will more than double. (See *Figures 1-2*)

Yet, they don't foresee a comparable increase in results. When asked what they expect in terms of a realistic increase in alpha throughput, most (47 percent) said a 25 percent increase, while the others were split on either side of that:

¹ The survey was not a random sample and the data cannot be generalized.



Figures 1 and 2

Experiments that may show the potential for greater alpha are being passed over.

some optimistically expect a 50-percent increase and others a more modest 10-percent increase. This implies that at least some scenarios are not being run because of a lack of compute capacity. Experiments that may show the potential for greater alpha are being passed over.

Many funds already use cloud to some extent and seem willing to increase that use. Most of the respondents—58 percent—already use a mix of on-premises infrastructure and off-premises cloud computing to run their algorithms. Only 26 percent perform all analysis in-house. Asked to rank four options to increase analytics capacity and effectiveness, most picked “enabling use of cloud resources to improve capacity without moving data” as their top choice. “Adopting solutions that allow for a focus on processing power spend instead of infrastructure” came in second. The other two options: “Allow larger numbers of models and simulation to run on existing infrastructure” and “expand existing infrastructure to meet growing capacity and performance demands” tied for third place. (See Figure 3)

Clearly, respondents see cloud as a way to meet demand while keeping costs down. The attraction of turning what a high-capital expenditure—expanding on-premises infrastructure—into a more reasonable ongoing operating expense with elastic cloud capacity is strong. With cloud, they can turn up the algorithms, running models, simulations and back-tests in the cloud as needed, paying only for what they use. They can test ideas more quickly and cheaper than ever before.

- Allow larger numbers of models and simulations to run on your existing infrastructure
- Expand existing infrastructure to meet growing capacity and performance demands
- Enabling use of cloud resources to improve capacity without moving data
- Adopting solutions that allow for a focus on processing power spend instead of infrastructure

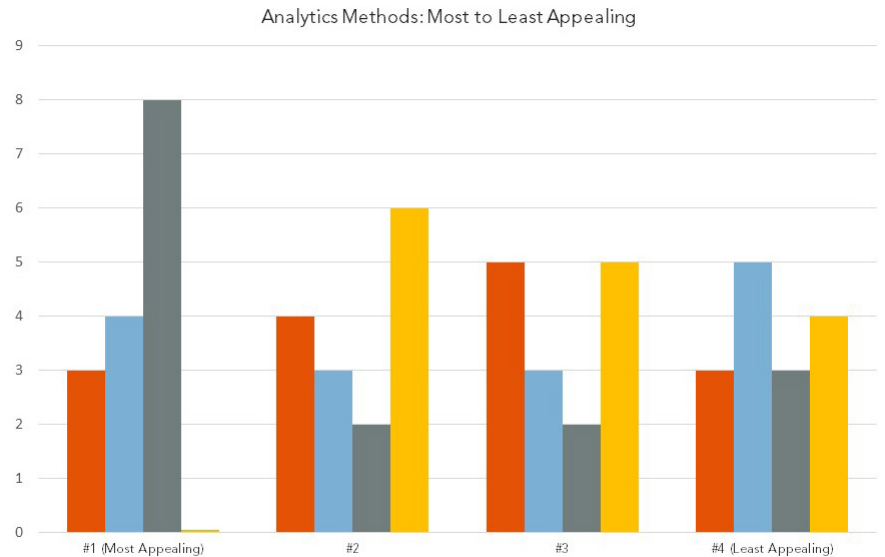


Figure 3

The preference of using cloud without moving data indicates one of two things: that some may still be leery about data living on the cloud, or that the process of moving the data is ideally avoided. In contrast, the desire to address processing power without spending on infrastructure signifies that the rental of cloud cores and storage is viewed as a viable alternative to data center expansion indicating that the security concerns once common are diminishing.

Balancing costs, speed, security and opportunity

IT departments in hedge funds and investment banks seek to address the quandary: How should it quickly provide for the spiky demand for processing power needed to factor data from unexpected events? After all, a fund's ability to gain alpha from these surprises often depends on the rate at which it can analyze the data.

IT managers feel frustrated by several factors that limit their effectiveness. Topping the list is lack of budget, followed by the "company boundaries" within which they must work, the time it takes to achieve results and the time it takes to evaluate new technologies. (See Figure 4)

The key is to attain this power and speed at the lowest possible cost, i.e. the best ROI. In the survey, 41 percent of respondents picked ROI as the measure that means the most when evaluating new products and services. The And cloud offers a clear ROI. For example, a company can spend over \$2.2 million

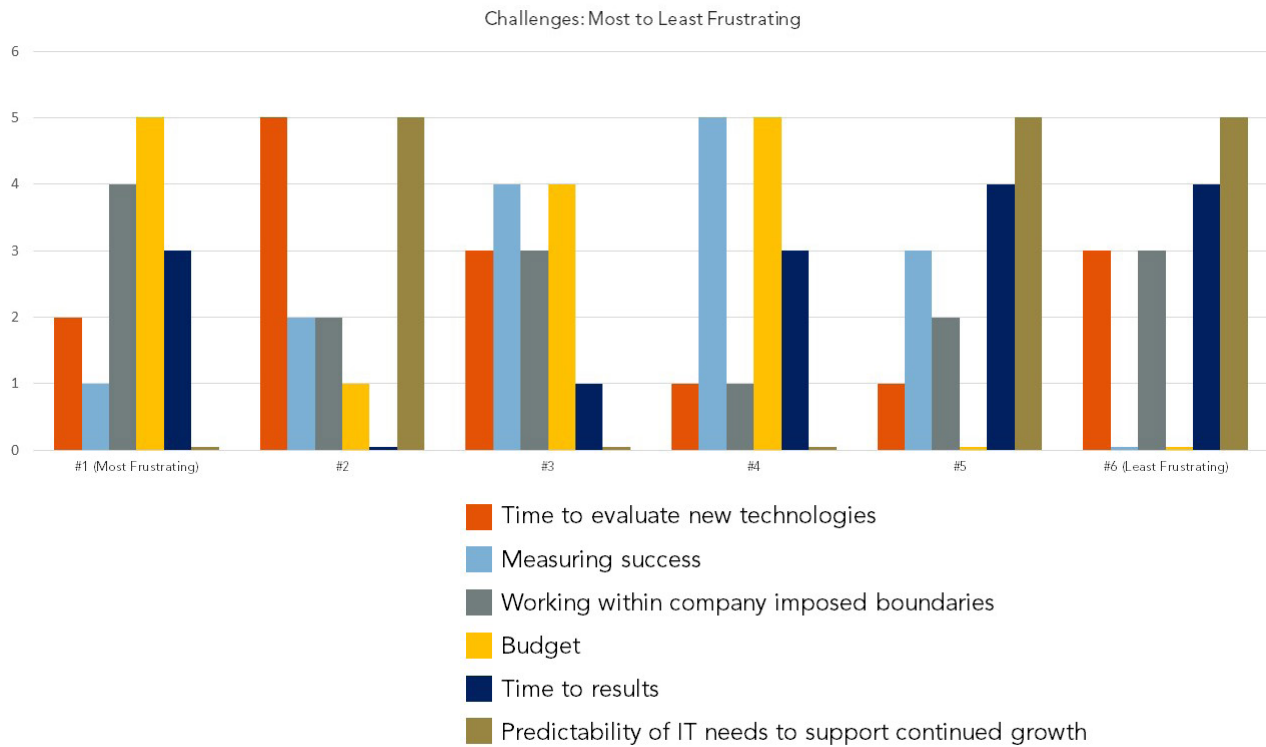


Figure 4

to add 5,000 cores to on-premises infrastructure or it can spend \$20,000 monthly to add 600,000 compute hours in the cloud. The second most important measure, picked by 35 percent of respondents, was the ability to more precisely predict risk, i.e. improving the quality of the analysis.

An often-overlooked factor that impacts both measures is latency. Whether adding compute cores to on-premises data centers or increasing core processing by cloud bursting, latency in fetching data slows the speed at which analysis can be performed, thus delaying time to results and potentially reducing alpha.

There are different types of latency that can limit alpha throughput in on-premises infrastructures and in hybrid scenarios. In both scenarios, delays can occur among core processing clusters and between clusters and data sources due to insufficient bandwidth, or the ability of the storage infrastructure to handle the sheer number of requests for data that these parallel compute jobs demand. Avere Systems caching and tiering technologies address these latencies, enabling IT to balance on-premises and cloud computing depending on the needs for performance, cost and security. This helps IT meet hedge firms' need to run more simulations in less time and at reduced costs, while increasing alpha throughput and potentially lucrative investment scenarios.

A way forward

In summary, hedge fund IT managers are faced with a growing challenge: Meeting the need for more power and faster speed while justifying their spend by showing a good ROI. The more back-tests firms can run, the greater chance of proving the viability of an investment plan, thus increasing alpha. But huge capital investments in on-premises infrastructure may not be the most cost-effective way to meet short bursts of demand for processing power prompted by a world full of surprising events.

Cloud computing and hybrid architectures can be the solution, but only if performance is not compromised by latency issues. IT professionals that find ways to effectively manage latency and balance performance between on-premises and cloud computing can deliver solid performance that helps their firms prosper in this volatile environment.

These results come from an informal online survey conducted by an independent contractor of Avere Systems in the fall of 2016. The survey targeted senior IT professionals working in large hedge funds and investment banks in the United States and Europe. More survey results can be found at www.averesystems.com/FinServSurvey.

