

Letter to investors, Q4 2019

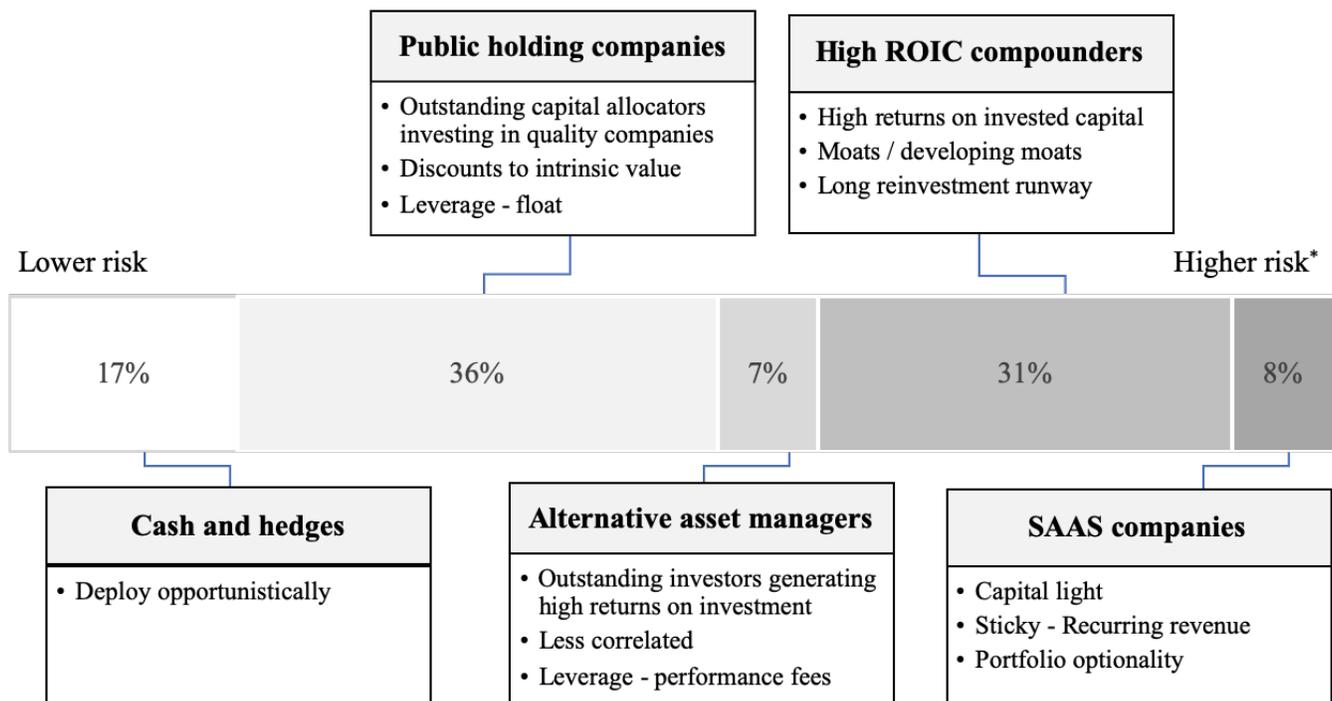
Performance

The Master Account¹, in which I am personally invested alongside SMA clients, returned 18.9% net in 2019, as reported by our fund administrator NAV Consulting. As of December 31st, 2019, the top ten positions comprised approximately 73% of the portfolio, and the portfolio held approximately 17% in cash.

I believe the portfolio is well positioned going into 2020, with a collection of quality businesses that have attractive valuations and optionality. I am actively looking to deploy cash, but with patience and discipline.

Portfolio composition

The goal is to compound wealth over the long term, responsibly, by investing in a) great public companies that have the power to endure, with long runways to grow through reinvesting cash flows at high rates of return, run by talented and aligned operators and b) the occasional special situation. The composition of the portfolio plays an important role in achieving in this goal. While I briefly mentioned my adoption of a barbell-type strategy in a previous letter, I wanted to expand on how I apply the concept in this letter. The diagram below shows how the portfolio is currently invested. The categories are somewhat subjective but align with the way I think about the various investment buckets.



As of December 31st, 2019

* Risk defined as a chance of permanent impairment of capital, and not volatility

¹ I had previously referred to this as the “model portfolio” but Master Account is more appropriate as this is a live account operating under real trading conditions vs a hypothetical model

With rare exception², the goal is to find companies with long runways and high ROIC opportunities. This holds whether the investment is in a company, in a collection of companies (through a holding company), with allocators of capital into great assets (financial management companies), or in earlier stage companies with the potential to develop future cash flows. Our current holdings have this attribute, in my opinion.

I look to strike a balance between risk of permanent capital impairment and the potential for upside gain. An ideal portfolio would have what theorists call “positive skewness” – it would participate in the upside but be protected from the downside, while still generating attractive returns. One way to achieve this might be to adopt what is called a “barbell” strategy (as popularized by financier and author Nicholas Nassim Taleb). The concept is to invest, say, 90% of one’s assets in ultra safe securities such as treasuries earning a risk-free rate of return. The remaining 10% would be invested in highly speculative instruments that have potential for large payouts, such as investments in biotech firms, startups or equity options. The problem with this approach is the majority of the portfolio does very little, and one is reliant on a very limited number of lottery type “bets” to generate any meaningful return, each of which has a low probability of being successful.

Borrowing loosely from this concept, in a far less extreme manner, I favor the idea of having the majority of the portfolio in robust businesses that collectively have, in my estimation, a low chance of experiencing permanent capital impairment while still having the potential to appreciate substantially over time. As I wrote in my first letter, if the US equity market has given investors a 6-8% gain per year over the last two centuries, then limiting one’s investments to higher quality companies with high ROIC should generate more attractive returns over long periods of time. Investments in well run diversified holding companies and high quality companies generating substantial cash flows are examples of this bucket. Sprinkling the portfolio with a few ideas that have the potential for high payouts can materially add to returns, and the risk is mitigated provided these positions are sized appropriately. In the event the more speculative investments do not pan out, over time the safer investments should make up for the loss. In the event the more speculative positions do well, the overall returns can be quite attractive. I will make the distinction here my choice of “more speculative investments” are still companies that I believe have a good chance of being successful, and not binary lottery type bets.

At the end of this letter I profile an investment in a smaller and higher risk software company that has the potential (though no guarantee) for dramatic upside optionality.

Finally, some thoughts on portfolio concentration. My personal preference is to have a 15-20 stock portfolio, which strikes a comfortable balance between investing with “conviction” and “being able to sleep well at night.” This is a manageable number of companies to follow closely, yet a large enough number to allow the portfolio to survive a mistake, provided it is handled appropriately (see the discussion on volatility in my Q2 letter).

To elaborate further. On the one hand, good investment ideas are hard to come by and take a lot of work to diligence. An investor should exploit these valuable opportunities when they appear by investing a large enough proportion of one’s portfolio to make a difference. Any position much smaller than 4-5% of the portfolio simply wouldn’t move the needle enough to really matter. (To be sure I have a few small positions in the portfolio, but these are a result of spin-offs where the starting position was much larger, due to fluctuations in share price after I bought shares, or because I am building a position with the intent to achieve a full allocation over time).

On the other hand, with a goal to “compound wealth over the long term, responsibly” I am acutely aware of the risks of over concentration. A more concentrated portfolio stands a greater risk of permanent capital impairment. To understand why, we can look at the principal of non-ergodicity. I’ll refer to a characteristically insightful presentation by Professor Sanjay Bakshi³ in October 2019 titled “Non-Ergodicity and its Implications for

² I am treading a fine line between articulating an approach that I can stick to with discipline and that investors can follow, and not boxing myself into an overly formulaic approach that might preclude me from investing in an attractive opportunity in an unknowable future

³ Professor Bakshi might not be a household name here, but he is a highly respected professor of behavioral finance and valuation at the Management Development Institute in India, and a successful practitioner

Businesses and Investors”⁴. He uses the following example of betting on coin flips: imagine the payouts are +50% for heads and -40% for tails. This would have a positive average expectancy of +5%. Nevertheless, an individual playing this game repeatedly is virtually guaranteed to lose over time. The reason is the difference between what is called the Ensemble perspective and the Time perspective. The Ensemble perspective is when the coin flips happen concurrently, say, 1,000 independent people flip the coin. The average of all those flips will be positive 5%, while within that average there will be many winners and losers. In contrast, the Time perspective is when the flips happen consecutively, say, one person flipping the coin 1,000 times in a row. After enough coin flips, the person will inevitably run into a streak of losers. Under the rules of this game, it only takes 4 consecutive losers to experience an 87% drawdown⁵, a level that is exceptionally difficult to recover from. In a game of coin flips, a 4-streak is surprisingly common.

Applying this to investments instead of coin flips, we can see that placing a large proportion of one’s capital into a single investment where the outcome is uncertain presents an unacceptable risk of eventually being taken out the game. This holds true whether the results are -40% and +50% (both entirely possible with stock returns) or a whole range of other percentages. The antidote is to spread one’s investments over enough concurrent opportunities to mitigate the consequence of any single investment going wrong, and to make the portfolio behave more like an Ensemble.

Discussion on a selection of portfolio positions with an average weight over 1%:

There was one exit this quarter and no additions. Positions with an average weight over 1% of the portfolio in the quarter were: Alliance Data Systems Corp., Burford Capital Ltd., CarMax Inc., CRH Medical Corp., Crossroads Systems Inc., Exor NV, Fairfax Financial Holdings Ltd., GCI Liberty Inc., Naspers Ltd., Oaktree Specialty Lending Co., Onesoft Solutions Inc., Prosus NV, Tencent Holding Ltd., Viamed Healthcare Inc., and Wandisco Plc.⁶

EXOR – Position trimmed

I sold some EXOR shares as the name had become over 28% of the portfolio, driven by the news of Fiat’s merger with Peugeot. This was purely a portfolio management exercise, as I continue to have strong conviction in the company for all the reasons mentioned in previous letters. From an operational perspective EXOR is executing. Its subsidiary CNHI, which receives a lot less press than Fiat, has announced a spin-off of its lower margin On Road business (Iveco). It will hold onto its Off Road business (tractors, farm equipment) which commands much stronger competitive positions, as evidenced by higher margins and returns on invested capital. It is clear CNHI, with the guidance of EXOR, will use a value creation strategy that is similar to the one that Fiat executed exceedingly well a few years ago as it spun out Ferrari and CNHI itself. While there is never any guarantee of success, I like to be aligned with operators with a strong value creation mindset. As for Fiat, it has agreed terms to merge with Peugeot⁷. This follows a well-publicized but ultimately unsuccessful attempt to merge with Renault, which was sunk by a mix of politics and high stakes corporate intrigue with Nissan (it is hard to miss the Carlos Goshn drama in the news). Though structured as a merger of equals, in comparing the relative share prices and dividend payments at the time of the announcement, Peugeot was effectively paying a 32% premium⁸. Hats off to John Elkann and team for both the quick pivot to Peugeot and a major step to realize the long held plan to improve the industry’s “structurally low and volatiles returns on capital” by rationalizing car platforms and attaining economies of scale.⁹

⁴ Non-ergodicity mostly developed by Ole Peters professor at London Math Lab. <http://lml.org.uk/people/ole-peters-2/>, Bakshi’s Presentation here: <http://bit.ly/2MOWbfJ> Jim Collins, author of Good to Great talks about it as well <https://fs.blog/jim-collins/>

⁵ $100\% - (1 - 40\%)^4 = 87\%$ drawdown

⁶ There is no assurance that any of the securities discussed herein will remain in an account’s portfolio at the time you receive this report or that securities sold have not been repurchased. It should not be assumed that any of the securities transactions or holdings discussed were or will prove to be profitable. See “Disclaimers” at the end for more details.

⁷ Fiat investor relations page <https://bit.ly/35TsnGc>

⁸ According to Jeffries analyst referenced in Reuters article <https://reut.rs/35UgHTE>

⁹ Sergio Marchionne’s notorious “Confessions of a Capitalist Junkie” presentation in 2015 <https://bit.ly/2TsqbTi>

Quarterhill – Exit

In my Q2 letter I profiled Quarterhill, a small technology holding company trading at a substantial discount to its components with a CEO who appeared to be driven to grow the company through a disciplined acquisition strategy and a focus on returns on invested capital. One of the risks that I outlined in my thesis was “Nothing happens – neither a deal is done (prices remain high, too many buyers looking to make SAAS deals), nor capital returned to shareholders. Stock becomes dead money.” In early Q4, I was disappointed to learn the CEO had voluntarily resigned. This surprised me as he had seemed very optimistic and had put in a lot of work to prepare the company to make acquisitions, including thoughtful compensation plans and a framework for post-acquisition integrations. I later learned that he decided to return to his previous employer, OpenText, where he had previously completed over \$3.8bn in acquisitions, benefiting from the larger company’s scale and resources, and the ability to focus on M&A without the added responsibilities of managing shareholder expectations.

Quarterhill did not appear to have a robust plan to find a suitable replacement, and so I promptly exited the position. Fortunately, the shares had been acquired with enough margin of safety that our portfolio was not negatively impacted.

WANdisco Plc

WANdisco is a small cap company traded on the UK AIM market. It has a patent-protected method of ensuring huge data sets can be migrated from on premise servers to the cloud and kept continuously synchronized, with guaranteed consistency, availability, and no business disruption. Something that was very hard to do at scale up until now. This is important for big data analytics, where enterprises seek to mine vast troves of unstructured data to generate actionable business insights (e.g. recommendation engines – where online retailers make purchase suggestions based on a user’s browsing history).

While the technology has been validated by some of the largest cloud vendors (e.g. Microsoft, Amazon, and IBM), for a number of reasons, the company has to date been less successful in ramping revenues. There is a very good opportunity for that to change following Microsoft’s decision to “white label” WANdisco’s product and sell it under its Azure label. Furthermore, WANdisco has formed a partnership with Databricks, a cloud-based data analytics company on the Azure platform with a \$6.2bn private market valuation and an impressive investor base (including Microsoft).

It is noteworthy that within 2 years of Databricks itself being white labeled on Azure in 2017, revenues ramped from approximately \$30m to \$200m and 5,000 customers. Importantly, those 5,000 customers to date have only uploaded a small proportion of their data to the Azure Databricks cloud, given the aforementioned difficulties of migrating huge datasets and keeping them in sync with on-premise servers. I believe both Microsoft and Databricks are incentivized to promote WANdisco’s data migration solution for their own commercial interest (more data uploads = more analytics revenue for Databricks and more storage/compute revenue for Azure). This could precipitate a significant revenue ramp for WANdisco. To put the opportunity into context, WANdisco has a market capitalization of approximately £215m with run-rate revenues of approximately £36m. It has a recurring royalty revenue model with 95%+ gross margins, and a directly addressable opportunity of at least \$1.5bn through the above-mentioned deals. I have included a detailed writeup at the end of this letter.

Closing thoughts

I wish everyone health and success for 2020 and look forward to updating you in the Q1 letter.

Samer Hakoura
Alphyn Capital Management, LLC
January 2020

WANdisco PLC (WAND:LN) – as of December 31st, 2019

WAND:LN listed on AIM UK

Share price: £4.45

Shares out: 48.2m

Market Cap: £215m

Net Cash: £12.8m (As of June 2019, £16.5m raised subsequently in November)

Infrastructure-As-A-Service play with significant recurring revenue ramp in 2020, following the recent signing of a co-engineered first-party integrated cloud solution with Microsoft Azure.

Thesis Summary

WANdisco is set for meaningful revenue inflection in 2020 due to the recent signing of a first-party integrated engineered cloud solution with Microsoft Azure, combined with partnerships with Databricks and Neudesic who are key cloud vendors on the Azure platform.

Databricks is a provider of cloud-based analytics with a \$6bn private market value, and investments from a who's who list of VCs as well as Microsoft. Databricks increased revenue by a factor of 6 and gained 5,000 customers within two years of establishing Azure Databricks as a first-party native service on Microsoft's cloud. It achieved this with access to only a small fraction of customer data, given the difficulties of migrating large mission-critical live data sets to the cloud.

WANdisco is an Infrastructure-As-A-Service software company with a patent-protected method to migrate very large live datasets and keep geographically dispersed servers continuously synchronized, with guaranteed consistency, availability, and no business disruption. Microsoft has now co-engineered a first-party integration with WANdisco and is sponsoring a joint Azure-Databricks-WANdisco-Neudesic campaign directed at approximately 1,000 customers specifically to promote migration and cloud analytics.

The potential implications for WANdisco are significant. Run-rate revenues are approximately £36m, with a credible path to over £100m within the next 2 years and £300m+ beyond that, at which point the company will generate 50% operating margins. As revenue inflects, WANdisco will benefit from impressive operational leverage given its 90%+ gross margins and largely fixed costs due to a royalty-based recurring revenue business model.

Overall, WANdisco has a 4,900 petabyte, \$1.5bn Annual Recurring Revenue opportunity, and this is not priced into the stock which only has a ~£200m market capitalization.

Cloud is the defacto standard for analytics, development and compute

Broadly speaking, the cloud market is very large and growing rapidly. IDC estimate that by 2022, 80% of cloud services will be concentrated in the big 4. And cloud services in general have become important sources of revenue. Q3 2019 "cloud" revenues were: Amazon - \$9bn, growing at 35% year-on-year, Microsoft - \$11bn growing 27%, Google (in Q2) - \$8bn with plans to triple its sales force "over the next few years," and IBM - \$5.3bn, growing 8%.

Over the last decade enterprises have become interested in the potential for Big Data analytics. The idea is that the massive amounts of unstructured data generated by customer clickstream data, tweets, videos, pictures etc., could be mined to generate actionable business insights. The advent of machine learning and AI has enhanced the appeal by promising predictive analytics, for example: recommendation engines – where online retailers make suggestions based on a user's browsing history, churn analysis – finding common patterns in the behavior of customers who are about to leave a service, and intrusion detection – monitoring networks and systems for malicious activity.

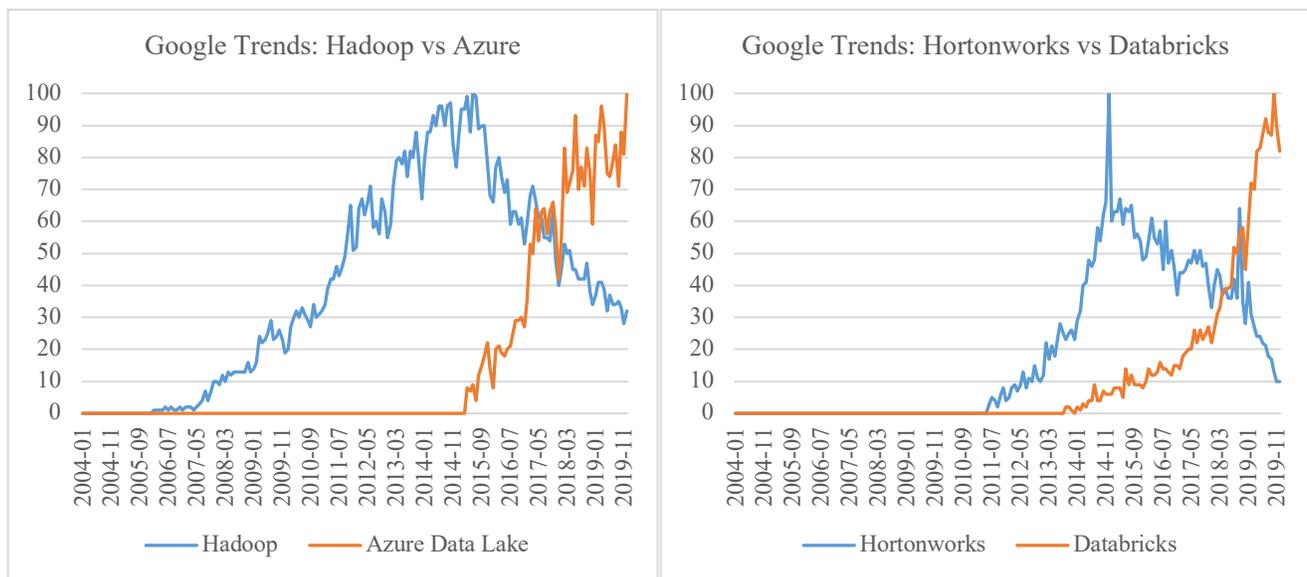
The amount of data required for these functions is massive. The data is frequently stored in data lakes, data repositories that stores large and varied sets of raw data in its native format. Traditional relational databases, which were developed in the 1980's, are not particularly well suited to hold this data. They require data to be organized

according to strict predefined schema, which does not suit the “unstructured” and dynamic nature of social media data, for example. The tabular nature of traditional databases also makes them less efficient and less easy to scale to handle massive quantities of data.

As a result, alternative data systems were developed, such as Hadoop, an open source software framework to store and processing data over distributed clusters of cheap/commoditized hardware. Companies made significant investments into their core on-premise Big Data infrastructure.

The level of interest in Hadoop peaked in 2015 and has been declining ever since, as the cloud has instead become the de facto development and deployment environment. With on-premise Hadoop companies had to budget for complex deployments, which were disadvantaged in terms of their speed, reliability, flexibility, ease of use, and cost. In contrast, cloud vendors could offer better performance, scalability, reliability, availability, a diverse set of analytic engines, and massive economies of scale.

The Google Trends charts below show the decline in interest in Hadoop and concurrent rise in interest in Azure Data Lake. A similar pattern is seen with interest in Hortonworks (one of the top distributors for Hadoop which was recently forced to merge with chief competitor Cloudera due to diminished business prospects) vs Databricks which provides cloud native analytics.



Source: Google Trends

To avail themselves of the benefits of the cloud, companies need a way to migrate their data from on-premise to the cloud as seamlessly as possible, without downtime and with no data loss. The migration opportunity is large. WANdisco estimates its directly addressable market is 4-6Exabytes of data amounting to c\$1bn - \$1.5bn of ARR.

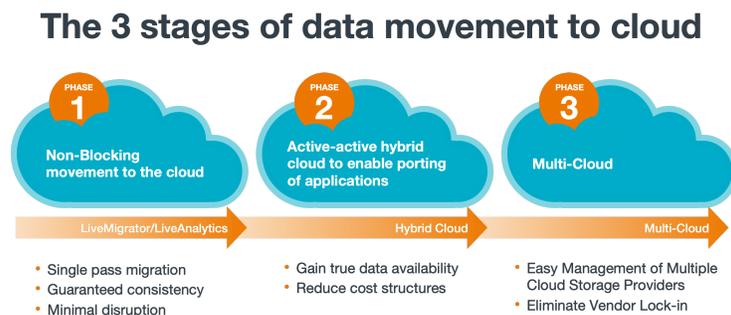
In an August 9th, 2019 note Stifel gave some further estimates of the broader migration Total Addressable Market:

- Data Center to Data Center: approx. \$4bn TAM
- Data Center to Cloud: approx. \$0.5bn TAM
- Physical Data Movement (e.g. Amazon Snowball): approx. \$10bn TAM
- Cloud to Cloud: approx. \$13bn TAM
- Total: approx. \$28bn TAM

WANdisco's solution

This shift from on-premise Hadoop to cloud (and hybrid-cloud) is very relevant to WANdisco. Large data sets are difficult to move, while enterprises have high expectations in terms of data availability, consistency, security, and auditability. A typical enterprise might have 3 petabytes of data, which would take a significant amount of time to migrate, during which time any newly generated data would be out of date. For reference, AWS Snowball's website states it would take 120 days to transfer 100 terabytes of data (1/10th of a petabyte) to their web servers over a 100Mbps internet connection, or 12 days over a 1,000Mbps connection. Their alternative solution for large (petabyte scale) data transfer is AWS Snowball, which entails physically delivering a ruggedized device with tamper evident locks, having the client copy data and return by courier for upload, while managing chain of custody logistics; a process that takes approximately a week and makes syncing live data a challenge.

While some competitors have "near-zero" downtime performance, which works for archive-quality data, WANdisco's is the only solution that works for continually changing live datasets. WANdisco's Fusion platform is agnostic as to storage environment and allows for data to be continually synced during the entire migration phase, with guaranteed consistency, availability and no business disruption. Furthermore, once migrated to the cloud, it gives customers the ability to continuously replicate data between on-premise and cloud, or indeed between multi-cloud vendors, at vast scale.



Source: WANdisco H1 2019 presentation

WANdisco holds 21 issued and 25 pending patents for its technology solution, called Distributed Coordinated Engine (DConE). The technology has been validated through partnerships with Microsoft Azure, Amazon Web Services, Google Cloud, Oracle, and others as well as OEM relationships with IBM and Alibaba. It was developed over five years by technical co-founder Dr Yeturu Aahlad, who was a distributed systems architect at Sun Microsystems. DConE allows multiple instances of the same application to operate on independent hardware, where a change on one server in one location is immediately replicated on the other servers. It is based on a mathematical algorithm, called Paxos, which is designed to keep a network of computers in sync given that one or more computers may occasionally fail and connections between computers are occasionally unreliable resulting in message delays/loss. In contrast with other solutions which typically have a master-slave configuration and present a potential single point of failure (data is lost if the master server goes down), Paxos treats all computers on a distributed network as peers, and obtains "consensus," where any change to data at a single computer on the network has to be agreed on by a majority of its peer computers, before being accepted by the system and then synchronized amongst all computers.

While Paxos works on computers in a LAN (local area network), DConE works over Wide Area Networks (WAN), where servers can be thousands of miles apart, as found in practice on the cloud or over distributed enterprise servers. It achieves this scale and efficiency by reducing the number of peers required to form a quorum to achieve consensus, amongst other enhancements.

For use case examples, WANdisco provides the examples of AMD Inc. using Fusion to provide continuous replication of critical business data to the Azure cloud for disaster recovery, a "Very large multi-channel retailer" using Fusion to synchronize data between disparate retail locations and central systems to use for predictive

customer analytics, and Daimler using Fusion to synchronize development and autonomous vehicle data to the cloud for cost effective compute beyond their in-house capacity.

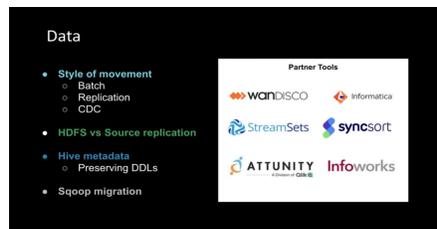
Competition:

There are three alternative solutions for data migration at scale. The first is point-in-time (vs continuous) static copy, with the obvious disadvantage of resulting data inconsistency between source and target data repositories. It is a manual and labor-intensive process to reconcile differences, often requiring the development of custom code. This is the method that the Hadoop distributors Hortonworks and Cloudera employ.

An alternative solution is Change Data Capture (CDC), provided by companies such as Attunity, Striim, and Streamsets. As the name implies, CDC is a process of capturing changes made at the data source and copying the changes to a destination. Some limitations are that the process is uni-directional, data is replicated one way from source to target and requires one location to be the source of “truth”, and the process requires vendor-specific APIs (e.g. from Amazon S3 using their API to the target API). The process is suitable for archive quality data, whereas Fusion achieves true 2-way active/active replication across multiple peers and is “non-blocking,” which means companies can continue to work on their on-premise data while it is being replicated.

Finally, there is Dual Ingest, as used by Google Cloud Platform. Data is ingested through a “load balancer” which copies the data to two locations. Limitations include: any disruption at either target servers causes the data to diverge, which means the process requires constant attention and administration.

In a November webinar on Hadoop-to-cloud migration, a senior Databricks representative (more on the Databricks partnership below) commented “*if you needed an absolute live replication solution, you could use WANdisco ... it just works.*”



Source: Databricks webinar

The opportunity – why now?

Over the course of 2019, WANdisco has announced a series of partnerships with Microsoft, Databricks and Neudesic, that provides a compelling use case for customers which I believe will fundamentally change the revenue trajectory of the company.

WANdisco Fusion is now a deeply embedded co-engineered first-party service on Microsoft Azure. In plain English this means that Fusion will be white labeled and fully integrated on the Azure platform, appearing like a native Azure feature, enabled with a single click and including integrated billing. This greatly reduces required planning and simplifies the migration process for customers.

Azure Databricks, also a first-party service, provides a single platform for data analytics and machine learning on the Azure cloud. Its greatest strengths are a zero-management cloud solution and an easy to use, interactive environment for data scientists, engineers, and business analysts to collaborate. It was founded by the original creators of Apache Spark, a more modern open source data processing engine, that can efficiently handle petabytes of data across very large clusters of computers, which makes it particularly well suited for large scale data processing; for example it is 100x faster than MapReduce, the older framework for Hadoop. The company has a global customer list, including Shell, Viacom and HP. It has raised a cumulative \$900m in venture funding (at a

\$6.2bn valuation) from the likes of Andreessen Horowitz, New Enterprise Associates (NEA), Tiger Global Management, BlackRock, and Microsoft itself.

Neudesic is a System Integrator that is Azure certified and a closely aligned Databricks partner. Moreover, Neudesic was given the Microsoft 2019 MSUS Partner Award for Intelligent Cloud – AI and Machine Learning.

The ensemble provides a compelling value proposition for customers: use WANdisco Fusion to seamlessly move massive amounts of data to the Azure cloud to run advanced analytics with Databricks, enabled by Neudesic.

To my knowledge there are only 3 first-party arrangements with Azure: Databricks, NetApp and now WANdisco. Brendon Howe, NetApp's GM of Cloud Storage Services, has repeatedly stated that he expects "revenue ramp rate will be significant." (Goldman conference February 18th, 2019).

Of particular note, Databricks achieved significant customer adoption and revenue acceleration since it was integrated as a first-party Azure service in November 2017. Databricks gained 2,000 customers and achieved \$100m in run-rate revenues in one year following the integration, and 5,000 customers and over \$200m in run rate revenues in year two. As I understand it, the majority of these customers are currently only running a very small subset of their data in the Azure Databricks cloud, and both companies are therefore incentivized to promote WANdisco's migration technology to increase data volumes for their own commercial interests. Microsoft have sponsored a 5 terabyte for 5 days migration/analytics campaign that is currently being marketed to approximately 1,000 Databricks target accounts.

Why is this not priced in?

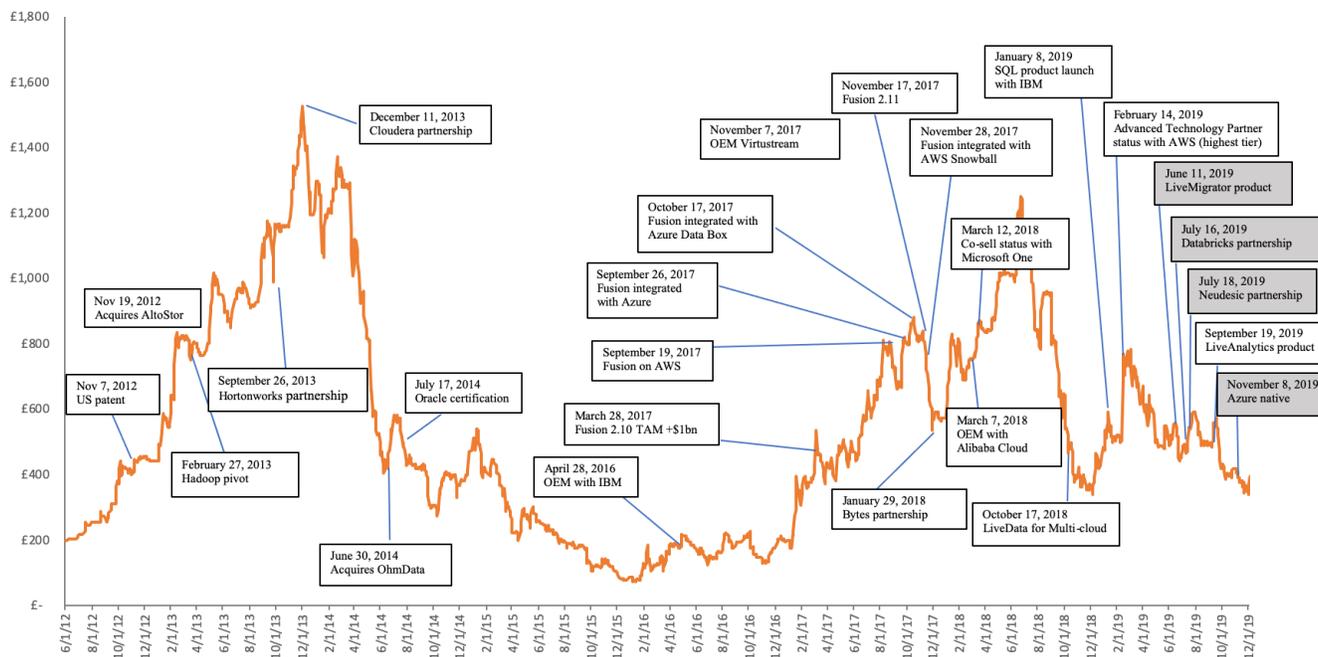
I believe the staggered manner in which announcements were made are part of the reason. WANdisco announced its LiveMigrator product to focus on solving the migration issue in June 2019. This was followed in July with the Databricks and Neudesic announcements and a mention that the company was working on an integration with an as yet unnamed partner. It was not until November 8th that Azure was announced as the major partner. In fairness it took Databricks a year to complete the integration, and NetApps have talked about "The step of making these as native first-party services is much more complex ... probably didn't fully scope the amount of effort to accommodate scale in steady-state that's required to be a truly integrated service" adding "it's a pretty lengthy journey and a fairly significant effort to do it with the goal of it being as natural as all the others. It's a pretty high bar."

The November 8th announcement contained the line that the partnership was only expected to become revenue generating in full year 2020 and was followed by a 4% share price decline. This only added to the perception of WANdisco being a "show me" stock with a history of big announcements that yielded only modest revenues followed by strategy pivots.

- In 2012 WANdisco's solution was initially aimed at the software development industry market, valued at approximately \$800m in 2012. HP was an early customer, with 50,000 users developing code working across 35 data centers. However, a competing open source distributed version control system called Git established itself as the dominant platform.
- In 2013 WANdisco pivoted to the Hadoop market with "NonStop NameNode" and signed partnerships with Hortonworks and Cloudera, at the time the two largest Hadoop distributors. While powerful, the software was fairly invasive and had to be implemented with great care to ensure replication worked correctly. This was an impediment to roll-out.
- In March 2015, WANdisco revamped its solution into the Fusion platform which was easier to deploy. Unfortunately, this coincided with the aforementioned peak Hadoop.
- Throughout 2017 to 2018, WANdisco worked on certifications with the major cloud vendors, which lead to some modest sales increases.

- Finally, in 2019, WANdisco refocused on the specific opportunity of migration to the cloud, via a go-to-market partnership with Databricks that has already demonstrated the potential for revenue acceleration achievable through a first-party integration with Microsoft Azure.

Share price developments annotated with significant strategic developments



Source: Company Regulatory News Service announcements

Optionality in the numbers

WANdisco is positioned for a significant revenue ramp as the new partnership starts to bear fruit. The sell-side analyst community, while acknowledging WANdisco’s significant recent strategic developments, have not yet updated their models, preferring to avoid any “forecast risk.” This sets the company up for material gains as it outperforms near term expectations (Munger’s stock market as a pari mutuel system where investors make the most gains from mis-priced bets).

(US\$ m)	H1 2019	H2 2019E	Medium-term	Longer term	Analyst 2020 Estimates	
	Actual	Pro-forma ann.	Estimates	Estimates	Stifel	Edison
Revenue	\$6.0	\$36	\$100	\$300	\$26.9	\$40.7
Gross Margin	93.7%	94%	90%	90%	91%	90%
SG&A	\$22.1	\$45	\$60	\$100	\$42.7	\$37.43
Operating Profit	-\$16.5	-\$11	\$30	\$170	-\$18.2	-\$0.8

Source: Alphy Capital Management, equity analysts, WANdisco announcements

In its H1 2019 interim results, published on September 25th, WANdisco was “confident in FY19 revenue guidance of \$24M.” This is comprised of \$6m H1 revenues, \$15m in known H2 late-stage deals and \$3m in H2 partner-driven sales. Annualizing the implied \$18m H2 revenue number yields a \$36m run rate revenue. These numbers do not include a potential additional \$14-17m in “other 2019 pipeline,” and do not include the Azure-Databricks-Neudesic deal, which is “expected to become revenue generating in full year 2020” (as per November 8th announcement).

I believe the path from here to \$100m is credible. A typical large enterprise has approximately 3 petabytes of data. At a price of approximately \$250k per petabyte per year for migration and then continuous replication, this translates to approximately \$750k in Annual Recurring Revenue potential per customer. Therefore, WANdisco would need only 133 customers to generate \$100m in revenues, significantly fewer than Databricks' 5,000 customers.

To sense check this: WANdisco views its directly addressable TAM for migration to be “in the region of four to six exabytes of data, amounting to between \$1.0 billion and \$1.5 billion in potential revenue.” This implies a price of \$250k per petabyte. For comparison, in an August 8th, 2019 report Stifel estimated WANdisco's target market for migration to the cloud is 2,500 potential customers paying between \$200k and \$2.5m each (TAM range of \$500m to \$6.25bn). In the last 3 years WANdisco has signed contracts for between \$188k and \$1m per year for replication, migration or hybrid-cloud, via its partnerships with IBM, Microsoft and Amazon. \$750k in ARR falls within this range and seems plausible. Moreover, there is potential for significant revenue expansion. For example, the September 25th, 2018 data migration and hybrid-cloud contract with Daimler (via Microsoft partnership) was for \$200k/year and represented only 3% of the customer's data pool. This implies a potential revenue opportunity of over \$6m from one large customer.

WANdisco will enjoy significant operational leverage and operating margin expansion as revenues scale, as per the table above. Recurring royalty payments from channel partners using their own sales forces and marketing functions reduces COGS and minimizes onboarding overhead. WANdisco's SG&A costs are mostly fixed, composed of staff and amortization of capitalized R&D costs.

Valuation

When it comes to valuing emerging SAAS companies, my strong preference is to try and be directionally right and not impose false precision. Should the stock once again capture investors' collective imagination, business could command a valuation in excess of 10x revenue. Databricks, with run rate revenues of \$200m recently raised \$400m at a \$6.2bn valuation.

The company currently trades at 6x its run-rate revenues of \$36m (please note currency – market cap in £, revenues in \$, ignoring the net cash balance). For reference, in Qlik acquired Attunity in February 2019 for \$560 million, 6.5x 2018 revenues of \$86m. Rounding down to 5x the medium term \$100m revenues, and assuming 48m shares outstanding, and a 0.76 \$ to £ conversion implies a price of approximately £8 in the next 2 years. Alternatively, basing a valuation on normalized 15x multiples on \$30m in medium term pro-forma operating profits results in a £7 share price. There remains the potential for significant further upside should WANdisco execute effectively on its longer-term migration opportunity and then on the ability to provide continuing hybrid and multi-cloud replication.

Aligned management team and investor-friendly financings

Executive insiders collectively own approximately 11.71% of the company and are deeply aligned with shareholders. WANdisco has a core group of high-profile investors, including respected tech-savvy funds and prominent family offices (Bass and Davis), who have participated in equity financings done at significant premiums to the share price. My understanding is that the latest capital raise was initiated by shareholders in order to take any potential liquidity risk off the table.

- February 14, 2019 - \$17.5m raised at a 9.2% premium, approx. 5.85% of share capital
- November 25, 2019 - \$16.5m raised at a 23.2% premium, approx. 6.63% of share capital

Lastly, the sale of approximately 1m shares in association with the acquisition of Oppenheimer Funds by Invesco has put some temporary pressure on the stock price.

Investor	Holding	Notes
WANDisco directors	11.71%	
Invesco (acquired Oppenheimer Funds)	10.44%	Disposed of c1m shares
Davis Capital Partners	7.48%	Participated in raise
T Rowe Price	6.56%	
Swedbank Robur	6.22%	
Ruane, Cuniff & Goldfarb	5.78%	Participated in raise
Ross Creek Capital Mgt	3.49%	Participated in raise
Global Frontier Partners, LP	3.39%	Participated in raise. WANDisco NE Director
Capital Research Global Investors	3.31%	
FW GP Holdco LLC	2.94%	

Source: Discussions with company and RNS filings

Catalysts and Risks

Catalyst: updates on Microsoft - Databricks deal, customer wins and revenue ramp.

Risks: quite simply the company now has to execute. The company should soon deploy its first petabyte-scale migrations, and these need to run technically without a hitch under live commercial conditions. The software has been reworked, integrations completed, and partnerships struck. Should the revenue ramp not materialize or prove to be anemic as in prior years with previous iterations of product/partnerships, this would be a fundamental impairment to my thesis. We are waiting to hear how the 5-terabyte marketing campaign performs and the extent to which customers then decide to upload significantly more data. Furthermore, we need to hear finalized pricing – that the \$250k-\$360k per petabyte assumption is realistic.

Disclaimer

The description herein of the approach of Alphyn Capital Management, LLC and the targeted characteristics of its strategies and investments is based on current expectations and should not be considered definitive or a guarantee that the approaches, strategies, and investment portfolio will, in fact, possess these characteristics.

Reference or comparison to an index does not imply that the portfolio will be constructed in the same way as the index or achieve returns, volatility, or other results similar to the index. Unlike indices, the model portfolio will be actively managed and may include substantially fewer and different securities than those comprising each index. Results for the model portfolio as compared to the performance of the Standard & Poor's 500 Index (the "S&P 500") for informational purposes only. The S&P 500 is an unmanaged market capitalization-weighted index of 500 common stocks chosen for market size, liquidity, and industry group representation to represent U.S. equity performance. The investment program does not mirror this index and the volatility may be materially different than the volatility of the S&P 500.

Performance results of the master portfolio are presented for information purposes only and reflect the impact that material economic and market factors had on the manager's decision-making process. No representation is being made that any investor or portfolio will or is likely to achieve profits or losses similar to those shown.

Results are net of all standard fees calculated at the highest rate charged, expenses and estimated incentive allocation. Model portfolio returns are inclusive of the reinvestment of dividends and other earnings, including income from new issues. The return is based on annual returns since inception and does not give effect to high water marks, if any. Returns may vary for investors who are restricted from participating in new issues.

Hypothetical performance results are unaudited and do not reflect actual results of any accounts managed by Alphyn Capital Management, LLC. Hypothetical performance results are for illustrative purposes only and are not necessarily indicative of performance that would have been actually achieved if an investment utilized the strategy during the relevant periods, nor are these simulations necessarily indicative of future performance of the strategy. Inherent limitations of hypothetical performance may include: 1) hypothetical results are generally prepared with the benefit of hindsight; 2) hypothetical results do not represent the impact that material economic and market factors might have on an investment adviser's decision-making process if the adviser were actually managing client money; 3) there are numerous factors related to the markets in general, many of which cannot be fully accounted for in the preparation of hypothetical performance results and all of which may adversely affect actual investment results.

There is no assurance that any of the securities discussed herein will remain in an account's portfolio at the time you receive this report or that securities sold have not been repurchased. The securities discussed do not represent an account's entire portfolio and, in the aggregate, may represent only a small percentage of an account's portfolio holdings. It should not be assumed that any of the securities transactions or holdings discussed were or will prove to be profitable or that the investment recommendations or decisions we make in the future will be profitable or will equal the investment performance of the securities discussed herein.

The graphs, charts and other visual aids are provided for informational purposes only. None of these graphs, charts or visual aids can and of themselves be used to make investment decisions. No representation is made that these will assist any person in making investment decisions and no graph, chart or other visual aid can capture all factors and variables required in making such decisions.

This report is for informational purposes only and should not be construed as investment advice. It is not a recommendation of, or an offer to sell or solicitation of an offer to buy, any particular security, strategy or investment product. Our research for this report is based on current public information that we consider reliable, but we do not represent that the research or the report is accurate or complete, and it should not be relied on as such. Our views and opinions expressed in this report are current as of the date of this report and are subject to change. Any reproduction or other distribution of this material in whole or in part without the prior written consent of Alphyn Capital Management, LLC is prohibited.