US Research Tour 2016



Table of Contents

Introduction
Executive Summary
Trends6
Fast forward 20177
Data8
The Innovative Enterprise11
Security15
Actifio
Airware
Arundo 23
Black Duck
CaféX
Cazena 31
Crossdeck
Databricks
Datameer
DeltaID
Exabeam
GigaSpaces
Gigster
Gigya51
GoodData
Joyent

US Research Tour 2016



Lookout
MarkLogic
Menlo Security
Mesosphere
Microsoft
Oblong75
Okta77
Orbital Insight
Origami Logic
PinDrop
Perlert
Skysafe
Splunk
Springpath inc
SS8
Stack Overflow 101
Striim
Synapse.ai
Talena



Introduction

This was our 20th annual Research Tour for Linda and I and our second as Ntegra Greenside. It was undoubtedly one of the best. It is hard for us to believe that the Tour has survived and prospered for so long and that it goes from strength to strength. Back in 1996, we managed to see around a dozen companies during the week but this year it had grown to almost 40, many of which were mind-blowing.

To be able to see so many remarkable companies in such a short period of time is both stimulating and challenging. The speed of change continues to accelerate at a frightening rate and the Valley is driving disruption throughout the business world. All companies need to look at the Valley to see whence new competition might emerge.

The Tour is like a shot of adrenalin to both the delegates and to us. The feedback from the companies we saw about the quality of questions and understanding are a great credit to those who joined us on the Tour.

We would like to acknowledge the contributions made by all the organisations that helped us to put the Tour together and especially to our five host companies – Microsoft, Andreessen Horowitz, Splunk, Intel Capital and Sequoia Capital; to all those who presented to us; to those who helped us put the Tour together and most of all to the delegates who made it such a stimulating experience.

We eagerly look forward to next year's Tour.



Mike Portlock Chairman Ntegra Greenside



Executive Summary

In May 2016, a party of eighteen senior executives from the UK spent a week in Silicon Valley, California, visiting a selection of IT companies on the 20th Ntegra Greenside Annual Silicon Valley Research Tour.



The purpose of the tour was to learn about the latest developments in the deployment and application of IT, focussing on how these are being exploited by major organisations for business advantage. We received presentations 39 different from companies of varying size and maturity, all of which were actively involved in providing products and services to help organisations become more effective in their use of technology. In addition, we received presentations from our three host Venture Capital (VC) firms: Intel Capital, Andreessen Horowitz and Sequoia Capital.

Tech companies continue to attract unprecedented levels of investment. As in previous years, there were some strong recurrent themes across the companies we saw: mobile everything, the maturing of software as a service (SaaS) and enterprise readiness to adopt the Cloud. Whilst some of the themes remain consistent, it was the speed with which they are maturing that was truly impressive.

In the last few years we have seen the of rise the 'post Blackberry' Smartphone within the enterprise, with many Silicon Valley companies having a 'mobile first' mindset, a view which has transformed how we interact with a broad range of services both within and outside the enterprise. Smartphones have become the default end-user device as they are easy to use and can be used almost everywhere, thanks to ubiquitous high-speed wireless connectivity.

Whilst the Smartphone enabled enterprise remains important, we also see the increasing range of sensors, robots and other 'Internet of Things' (IoT) devices providing new and novel ways of interacting with services. It is likely that these will gain traction in the home first, but the blurring of our lives between work and home will ensure these new 'interfaces' impact on the enterprise as well. IoT was referenced by a number of our presenters, particularly as a catalyst for big data becoming even bigger.

The trend for demystifying big data continued this year, with many presenters talking about the democratisation of big data, whereby everybody can be a data analyst. The challenges of managing data were discussed, both in terms of storage and access. A number of the start-ups and more established companies have been concentrating on developing products, services and add-ons that enable business users to more easily 'wrangle' their data. A key enabler is provision of convenient, automated tools which translate, format, move data to and from the cloud, and generally ease the previously onerous task of managing huge volumes of information (even before the analytic queries have been run).

A similar paradigm exists for software. Our VCs predicted that in 2016 and beyond, we should expect a rise of designers ahead of coders; for the last ten years only coders built great software. New tools will enable everybody to be a great coder, so it is anticipated that designers will be king for the next ten years.

Data virtualisation and copy data management solutions are also becoming important because, although the old adage 'storage is cheap' still holds true, housekeeping multiple copies of large datasets is very expensive.

Applying a schema or plan to data when it is inserted into a database to enforce consistency is becoming less popular. 'Schema on Read' was one of the buzz-phrases of this year's tour. This innovative analytical strategy involves applying the plan or schema to data as it comes out of the database, rather than when it goes in. This approach allows for a much more versatile organisation of data and enables differing views of the same data to be (virtually) present.

UAVs (Unmanned Aerial Vehicles, or drones) featured on the tour for the first time in 2015. This year we saw further evidence of the growth and speed with which this market is maturing. UAVs have become yet another sensor platform providing huge volumes of data for driving analytics and insights. Al and machine learning advances mean new analytical techniques can be applied to this data, which also facilitates unlocking potential in historic data.

A great example of the innovative use of remote sensor data was Orbital

Ncegra Greenside

Insights, who have combined advanced analysis via a mix of AI, machine learning and data science to drive a range of applications exploiting digital imagery. For example, they use satellite images to measure car park occupancy at Walmart stores and use this information as a proxy for footfall and business performance. They can measure and report on this data, almost in real-time, way ahead of traditional approaches.

As in previous years, security was again a key topic. Threat landscapes are constantly changing and becoming ever more sophisticated. Prevention alone is futile, with hackers becoming more elusive and networks already compromised. Several start-up security firms asserted that enterprises should simply assume that they have already been breached, but they just don't know about it, yet! 72% of CEOs surveyed by KPMG in 2016 cited cyber risk as their top business risk and felt there was more work to be done to protect their organisation. Security by obscurity is no defence, but no one vendor can provide the protection that is required. A multi-layered, multivendor security architecture was deemed necessary. Enterprises need to invest in 'Red Teaming' and similar techniques to stress test their architecture and rehearse their countermeasures and responses.

Trends

This year's tour has reflected on some of the challenges that enterprises face

when adopting new technology, particularly from 'early to market' companies. We delved into security and big data in greater depth, given their seismic impact on the enterprise. What are the key messages or 'takeaways' for us this year?

- Enterprises need to have a clear strategy to work with emerging technology companies. Trying to depend on traditional governance, procurement and project processes will not work. The enterprise needs a 'Red Door' that shows companies the way in. Process, governance and culture need to be closely aligned.
- We should now regard the 'cloud' as safe for sensitive data and apps. This does not remove the need to ensure that solutions align with security architecture, end to end. However, saying no to a service simply because it is cloud-based could mean you lose out to your competitors.
- Finding and retaining staff will be as hard as it ever was. Think about building ecosystems for innovation and always be mindful that "you do not employ all of the smart people" (Mark Zawacki on our tour in 2014). Working more in the open (think 'open source for people') will be how successful companies grow and develop.
- Agility is critical for success. Enterprises need to increase focus on services based IT architectures, APIs and micro services.

 Operational excellence is as important as it always was. Cloud services mean that software asset management is increasingly important, as is data sprawl and sovereignty.

Fast forward 2017...

What are some of the trends that we expect will develop over the remainder of 2016 and through 2017?

- The disintermediation of data scientists from data science and coders from coding. Tools we saw this year, such as Datameer, are helping everybody become a great data analyst or scientist. Similar developments will happen in coding, where anybody will be able to write great code, so the value and demand for designers will increase. Increasing processing speeds coupled with infrastructure growth nearing 'critical mass' will accelerate the use of neural network based Machine Learning, Deep Learning and true Artificial Intelligence.
- Fast data is as important as Big Data. As enterprises more fully understand big data and start to use it to drive value, the velocity bar will be set ever higher. The enterprise must be able to rapidly process large datasets and analyse at pace to drive competitive action. Technologies such as Spark, Kafka and Streaming will see widespread adoption, driving near-real-time analytics, and fuelling the next

generation of big data growth.

- Open source will become standard for infrastructure software, with CIOs adopting an open-source, software-first approach. Some open source providers such MongoDB and DataStax (who we have seen on previous tours) have built monetization models whilst many others have millions of downloads, but no revenue. We expect more providers to find ways to monetize their offers.
- India is now the second-largest market in terms of connected users. 353 million, ahead of the US, and second only to China. This market will be strategically significant to US-based enterprise software companies. Cloud-based services are particularly appealing for high growth Indian companies, as they do not have the same investment in legacy infrastructure, thus simplifying implementation and operations management. This reflects mobile adoption in India, where rapid growth and lack of legacy enabled them to leapfrog a whole generation of technology (in this case wide scale deployment of fixed line networks). Bypassing entire generations of technology can be a major catalyst for achieving stellar growth.

Planning is already underway for our 2017 tour, with many of these themes and topics expected to play a major role. I hope to see you in 2017! Andy Jefferies



Data

Last year's research tour identified a number of interesting trends and developments associated with the management, manipulation and exploitation of huge volumes of corporate and social data. As predicted, the year proved to be an important one with, what had been considered just hype, becoming increasingly normal as organisations realised that data, big and small, is critical to ensuring best business outcomes. Throughout 2015/16 the continued, compound effect of Moore's Law coupled with improvements in compute, storage and network infrastructure capacity has enabled innovators to accelerate improvements in data processing techniques and capabilities. The convergence of these improved processing techniques and infrastructure capacity nearing 'critical mass' will permit new paradigms in Machine Learning, Deep Learning and Artificial Intelligence to become ubiquitous

This year, the companies we have seen have continued to drive rapid growth of systems that support non-relational or semi-/un-structured forms of data, as well as increasing their analytic capabilities at massive scale and speed. Established big data and analytics solutions are evolving and maturing to integrate well and align with enterprise IT standards. A number of the start-ups presenting this year have come about specifically to address the gaps in newer technologies that make them difficult to integrate with existing business systems and operating models. This trend is likely to persist as the rate of development of new raw capabilities increases. Big data startups are contributing, building and integrating components such as security, authentication, fine-grained role based authorisation and business continuity capabilities that customers expect from traditional enterprise management relational database systems. These, previously overlooked,

capabilities are now becoming key enablers and differentiators within the eco-system of emerging big data technologies, eliminating barriers to enterprise adoption.

As well as working to provide enterprise class systems, companies are configuring their solutions to enable both business users and data scientists to fully realise the value of their data assets. There is growing demand from business users to have the same self-service access to insights that they get from traditional data warehouse environments. Companies such as Datameer are blurring the lines between traditional BI/MI and big data, enabling business users to discover insights in any data via wizard-based iterative 'point & click' analytics and 'drag & drop' visualisations, regardless of the data type, size, or source.

Companies are rapidly negating the need for "swarms of experts in white lab-coats" to be continually nursing



corporate big data solutions. Selfservice, self-discovery and automated commentary (describing why insights and analytic results have come about) are what's expected by business users today.

Business users also want to reduce the time and complexity of preparing data for analysis, when dealing with a variety of data types and formats. We saw a number of companies, such as MarkLogic and GigaSpaces focusing on end user data preparation. Their customers dynamically can also asynchronously scale up or down the amount of storage and compute resources in the databases relative to the larger amounts of information stored in "data lakes". Storage of data is comparatively cheap compared to the cost of the compute resource needed to process it, so it makes perfect sense to use the elastic provision of resources in the cloud to ensure that compute is only paid for when it's actually being used. The effort to 'humanise' IT, enabling people to intuitively interact with systems, as opposed to systems asserting behaviours that enable interaction, is increasing. A number of companies demonstrated how data security can be moved into the background using biometrics and new algorithms to faster. enable non-intrusive authentication and authorisation.

Companies like GoodData are providing both the tools and the expertise needed for organisations to collect, analyse and exploit data, allowing users to quickly and easily see the impact of changes made on performance. This leads to a more human approach to performance improvement based on feedback loops and proven success.

Adoption of NoSQL technologies and a preference for storing data in unstructured schema-less form (where data is applied to a plan or schema as it's being pulled out of storage, rather than when it is written) were common themes. NoSQL and 'schema-on-read' databases are becoming an established part of the enterprise landscape as the benefits of schema-less database concepts become more recognised and understood. The Gartner Magic Quadrant for Operational Database Management Systems was previously dominated by the likes of Oracle, IBM, Microsoft and SAP. Now, however, NoSQL companies, like MarkLogic, outnumber the traditional database vendors in the leaders' quadrant.

Growth in the massively parallel processing (MPP) data warehouse segment has been slowing recently and the "death of the data warehouse" has been predicted. However, companies such as Cazena are driving a resurgence in the popularity and use of this technology in the cloud. Their solution provides self-service orchestration of cloud infrastructure in Amazon AWS RedShift and Microsoft Azure SQL Data Warehouse. Cazena uses these environments to provide, what it calls,



Data Mart as a Service (DMaaS) alongside Data Lake as a Service (DLaaS) configurations on Hadoop and other schema-less databases. This enables on demand provision of data processing and analytics platforms, with other technologies including Google BigQuery likely to be included soon, giving customers seamless access to best of breed workload engines and heterogeneous infrastructures which collectively provide the 'best of all worlds'. Something most enterprises would simply not contemplate on premise.

Apache Spark has rapidly become the big data platform of choice for an increasing number of large enterprises. DataBricks co-founder and Spark originator, Matei Zaharia, has seen it grow, from a small component of the Hadoop ecosystem, to provide significantly improved data processing speeds and become the largest current big data open source project. It will be interesting to see if Spark retains its "poster child" status in the months to come.

There was similar attention being paid to Kafka, a fast distributed publish and subscribe messaging mechanism. This is undoubtedly one to watch for next year, as it will impact the way new solutions can handle streaming data elastically with resilience in the cloud. It looks like a key enabler for managing and processing data from the IoT.

Data volumes from devices in the IoT are a further driver for petabyte scale

growth in the cloud. Established companies such as Google, Amazon Web Services and Microsoft are developing IoT services to enable data to move seamlessly to their cloud based analytics engines. Services that ease the pain of wrangling and conveniently storing this data, are enabling companies like Arundo to develop predictive solutions that raise asset utilisation and performance in industrial companies. They do this by combining sensor and transactional data with deep domain knowledge and experience to reduce maintenance costs and avoid unexpected outages using machine learning techniques.

Using machine learning algorithms to enable systems to learn how humans work (not the other way around) helps to spot patterns and connections between activities and performance, leading the way to more innate decision support and ultimately autonomous decision making capabilities. Driverless cars, drones and robotics pushing back the are boundaries of what's possible in this domain. Developing technologies so that interaction is driven by human preferences and needs, rather than technology capabilities, will come to the forefront during the next year. IT literacy may not be a differentiator for individuals soon, but human literacy and understanding will become an essential component of near-future systems.

Adrian Fern



The Innovative Enterprise

Almost every large enterprise has created a C-level post or responsibility for innovation. Recognition that innovation requires focus is a step forward. However, to what extent is the new breed of Chief Innovation Officers, Chief Digital Officers et al delivering? Why do many Enterprise Innovation agendas fail to meet expectations? Are expectations too high or do enterprises at scale struggle to innovate?

Typical factors constraining innovation are conflicting business goals and competing priorities. These are common challenges for larger organisations with broad functional, geographical and technical diversity. Governance and process frameworks can have a significant impact, for example; whether, where and how performance is measured: local vs. global indicators; or traditional budget / forecast financial cycles. Many of these features are perceived as necessary to control an organisation and its workforce, to understand return on investment, to understand risk and to comply with reporting requirements etc. Most are at odds with driving innovation, but this does not mean that large enterprises cannot be innovative.

The key to success

"If the highest aim of a captain were to preserve his ship he would keep it in port forever". Thomas Aquinas

To drive innovation in any organisation requires leadership and active engagement from the CEO and executive team.

Ground-up innovation by stealth does not work because organisations are

structured *not to* support change that evolves in this way. This innovation approach does not align with project processes, budget governance and / or security policies. In fact, most organisations are structured so that this form of innovation is actively discouraged or even targeted and destroyed.

The innovative CEO must encourage, or at times force, senior leaders to take a stake in the game and participate in new initiatives. Participants should be given the opportunity to explore new possibilities. Innovation should be explicitly linked to wider business goals such as growth or cost reduction. By drawing a strong linkage between business imperatives and innovation, the CEO can be certain that the company's innovation strategy is no longer abstract, but is directly connected to business performance. This linkage also supports integration with governance processes such as budgeting and planning, ensuring that appropriate financial and staff resources exist to enable change to be delivered.

One of the challenges faced by large organisations is how ready they are to innovate. Readiness should be the product of an honest assessment of process, capability (staff and technical), governance and, most importantly, culture. As Mark Zawacki from 650Labs said on our Research Tour in 2014, *"You don't employ all of the smart people".* A significant proportion of Silicon Valley innovation is based on open ecosystems, whether physical (in the form of facilities such as RocketSpace, the start-up accelerator run by Duncan Logan) or online communities such as Stack Overflow.

Agility, whether in process or tech, is also critical to innovation. We are seeing recognition that the Cloud is finally a safe place for the enterprise and has become the 'going in' position for most tech enablers. There is no doubt that successful cloud adoption drives agility and the ability to transform.

"No one today would propose a centralised, rigid, top-down organisational structure, where you cannot communicate across functions except through your bosses, but that was precisely what Alfred Sloan set up at General Motors, to great advantage at his time. With today's communications flexible and technologies, agile, creative networks make more sense and lead to much more productivity". Carlota Perez

Despite Carlota's comments made over a decade ago, how many large organisations remain structured along the lines of GM? How many have tried to implement enterprise collaboration technology such as Slack to break down silos, while process and governance continually reinforce?

Innovation needs to be integrated into an organisation, not bolted on.

Models for Enterprise Innovation

Internal R&D – In tech-centric, product based companies this type of innovation makes a lot of sense. It requires a long-term commitment, large budgets and is high risk. Pure innovation can create game-changing products/capabilities and associated profits; however, there are no guarantees. Apple generates around \$13 of profit for every dollar of R&D investment, Nokia towards the end of its life only made only \$2 of profit for every research dollar, highlighting how perilous running R&D is.

Acquisition – Buying a company to get your hands on their technology is fast, but potentially costly. Also, integration often presents a huge challenge - think Microsoft and Skype, Microsoft and Yammer, Microsoft and LinkedIn, Microsoft and anyone! The acquisition route is littered with examples of large companies that have failed to exploit the potential of their acquisition. Outside the world of technology companies, acquisition is less of an option to drive innovation.

Strategic Alliance – Partnering for 'winwin' outcomes sounds attractive, however, many marriages end in divorce or never get out of the 'prenuptial' phase. It is very easy to get buried in lengthy, hypothetical discussions regarding how the partnership will benefit all parties, rather than moving into action and experimenting with some live opportunities or problems. Our advice, if you seek the alliance route, is to implement a short-term framework to test the relationship on a real project then figure out how to build a longterm relationship.

Hire and retain the best talent – This is challenging and expensive for large enterprises. The best talent is often attracted to the early to market companies and predicting technology mixes that large enterprises should be recruiting for is notoriously difficult.

Red Door – This is a term used by Duncan Logan and, in our view, is critical to innovation within the Enterprise. Large Enterprises need to understand how to work with early to market companies. This covers everything from working spaces to contracts, to culture, to attitude, to intellectual property and to risk.

The Red Door

Start-ups enter here! $\rightarrow \rightarrow \rightarrow \rightarrow$



The Red Door is about attracting the early to market disruptors and enablers. The Red Door says "start-ups enter here".

Larger enterprises need to develop strategies that enable them to work in an agile way with the new high growth companies. This addresses a major concern of CEOs who feel that CIOs often do not move quickly enough to embrace new technology (77% of CEOs in KPMG's recent CEO survey [January 2016] support this position). Early to market companies are likely to be happy to work in the open and will shy away from lengthy discussions with inhouse legal teams. However, by building real partnerships (the enterprise gets first or early mover advantage, the start-up gets enterprise insight and real-world references), large enterprises can make genuine step changes.

P&G are a real example of how to do this. They launched a "Connect and Develop" program to systematically engage with partners outside the boundaries of the company to source new product ideas. This marshalled its famous process discipline to create a product innovation "Growth Factory".

In 2012, its head of technology announced that P&G had tripled its innovation success rate (the percentage of ideas that made it successfully to market as new products). Indeed, innovation has long been at the heart of P&G's success, which took its founding inspiration from Thomas Edison, who created the world's first industrial research lab. As former CEO, Bob McDonald, said: *"We* know from our history that while promotions may win quarters, innovation wins decades." [source Forbes].

How to work with Start-ups

There are many models for opening the Red Door and working with Start-ups. P&G have chosen to run their own accelerator; accelerators typically operate as a 'greenhouse' for ideas. In general, the idea and company exist, but it needs the right environment to grow. This may not mean moving the company to a new location; they may remain where they are, but they will gain access to a broad range of SMEs and mentors. An alternative model is to use an incubator, where initial, very early stage ideas are nurtured, typically in a co-working space.

The co-working space is critical as it not only provides cost-effective access to facilities; it also provides easy access to other founders and their expertise and knowledge.

Increasingly incubators and accelerators are running enterprise programmes where they help match enterprises with innovation challenges to potential technology enablers and disruptors. Ntegra's Innovation Lab operates on a similar principle; our model is to take a challenge from an Enterprise, identify the ecosystem of technology enablers that could address the problem and bring them together in an agile, cycle based Innovation Lab. This can be on or off-prem. We ensure we have NDAs and frameworks in place with the enablers, making it easier for the enterprise as they only have one contract to put in place; this one point of engagement effectively provides our clients with an Innovation Integrator. Creating chains of trust between the enterprise, accelerators / incubators, SaaS Apps and cloud providers is a key feature of being able to innovate successfully.

To be successful, innovation in the enterprise must not be limited to technology. Whichever model is adopted to drive innovation, business change must be included to ensure that the innovation can drive value and be successfully adopted by the organisation. A recent survey of CEOs by KPMG showed that many feel that in the next three years there will be more change than in the previous fifty. CEOs are very aware that how they embrace new technology will determine how effective they are, with 65% of respondents concerned that their business model could be disrupted by a entrant new [Source KPMG]. Governance models need to be developed to recognise the drive to be more innovative and agile, moving from monolithic projects to short sprints and continuous incremental change.

Andy Jefferies



Security

This year's tour identified a number of emerging trends in the IT security domain, with new companies highlighting a broader spectrum of concerns that should be considered holistically. Multiple layers of security aimed at managing specific risks are needed, and these should overlap to provide strength, depth and dependability. Ways of working are continually changing, leading to more opportunities for data to be targeted.

Organisations cannot assume they have enough security in place to prevent them from being compromised. They must continue to defend themselves but also assume they're about to be breached, and have appropriate controls in place to manage the aftermath of a security incident. The TalkTalk breach, last October, raised awareness about the possible scale and potential impact of breaches, particularly the media storm and reputational harm that followed. Unfortunately, when hackers stole identities from the online infidelity site, Ashley Madison, the effect was somewhat less galvanising. They had made much of the importance of security and anonymity, given the nature of what they were up to, but the almost comical nature of the breach. when it inevitably happened, did at least remind us that online privacy is very tenuous at best. It is clear that most organisations are still struggling to understand the risks they face and to take appropriate steps to be able to respond when they are compromised.

IT security is often described as an onion. Each time a layer is peeled back there is yet another layer to worry about underneath. Security teams are struggling to keep up with the new ways that individuals, good and bad, can access, use and interact with organisational data and information. New devices, new apps and new services all pose additional, cumulative risks and many do not have security controls built-in. Development of new services and apps by small companies on tight budgets, often mean security is poorly implemented, if thought of at all. This complex and diverse threat landscape is growing at an exponential rate and IT security teams simply do not have the time, resources or awareness to keep up with it.

In an attempt to mitigate, some organisations resort to using "red teams" to test their IT systems and provide the insights needed to update and improve their security controls (technology, people and process). Organisations can also create fake, network connected, machines (such as the TrapX honeypots) to divert hackers away, and enable security experts to analyse how the hackers actually go about penetrating these targets so that enhancements can be made to protect their *real* systems. These approaches help to provide some confidence and can be effective. However, as well as

N'tegra Greenside

protecting data, organisations need to ensure that the right people can access the right information, when they need it, often from external environments where there are many unknowns. These contradictory requirements have led to a number of big data and analytics companies, for example Prelert, to repurpose their capabilities and concentrate on individuals' usage patterns and behaviours. The ability to understand and baseline normal behaviour enables detection of abnormal activities and anomalies.

Hacker *journeys* are identifiable, in the same way that customer and employee *journeys* are, within time series data collected and logged across all the potential touch-points and routes through the corporate IT landscape. SS8 showed us how this can help to highlight risks from "suspects of interest" or "devices of interest", and illuminate areas where different security layers must interact and flex to form an appropriate defensive barrier which is only permeable to appropriate (normal) users and usage.

To harden and help improve the quality and lineage of new software products, start-ups such as BlackDuck and Tinfoil have developed automated scanning mechanisms that can be integrated into development pipeline processes. BlackDuck is able to detect open source components and libraries that are commonly included in builds, and hence provide developers or end-users with an assessment of their licence exposure and risk. Tinfoil scans newly compiled executables to identify all known security vulnerabilities and then raises defect reports in the developers' bug tracker (complete with instructions on how to debug the code and eliminate the vulnerability). Asserting Tinfoil into the Continuous Integration development cycle is like having an infallible security expert with coding skills in the development team.

Outside of key mission critical areas such as defence, in reality IT is currently largely uncontrolled due to the proliferation of uncontrolled devices and regular use of uncontrolled networks. Employees have little or no training and their knowledge of IT security is, at best, poor. They move organisations, between crosspollinating careless behaviours that lead to contagion in process vulnerabilities that, in turn, provide huge opportunities for hackers to exploit. For these reasons, security needs to be considered and planned at all levels: from employee identity; to building security into bespoke applications; to managing devices; to tracking intrusions.

The top, and most obvious vulnerability layer is no longer an individual's ID and password, but his or her identity itself. From just the name of an employee, a hacker can gain an entry point (for example, making use of social media to access personal details combined with a knowledge of an organisation's remote access solutions can often



provide enough detail to gain initial entry into corporate systems). Once "inside the lobby", brute force password attacks can be completed in less than a few second due to the immense computing power that is now available, on demand, in the cloud.

One innovative approach that reduces the chances of an employee making basic mistakes was demonstrated by Menlo Security. They protect against cyber-attacks from the web and e-mail by providing an 'Isolation Platform' which insulates content and eliminates malware in the cloud. Users' web sessions, and all active content, whether good or bad, is fully executed and contained within the Isolation Platform. Only safe, malware-free rendering information is delivered to the users' endpoint browsers. No active content leaves the platform, so malware has no path to reach an endpoint, and legitimate content does not need to be blocked in the interest of security.

This year, we've seen that already overstretched security teams are also having to think about, previously unforeseen threats and new technology responses from outside their traditional domain and skillsets. SkySafe provided an interesting description of their anti-drone systems which can be used to protect airspace and detect, and bring down drones that may be eavesdropping corporate secrets by listening in on conversations or photographing industrial or military

installations. Pindrop described their anti-fraud and authentication solutions for enterprise call centres. Within 10 seconds their technology can generate a risk score by simply "listening" to the background noise on a telephone connection and assessing whether the call is really coming from where it purports to be from. DeltaID made a good case for use of iris recognition as a better biometric than face and fingerprint. Face recognition is a challenging technology, with changes in lighting, hair style/colour and other looks resulting in a variable user experience. Fingerprints are a better and more popular biometric but they can still be problematic as they can be effected by weather, age, work, and many other factors. Using the iris as the biometric, results in a much better accuracy.

The above examples and many other interesting insights, from new companies, helped to illuminate the broad extent and continuing growth of the cyber-threat surface. This can be seen as a proxy for, and measured proportionally with, the rate at which new innovations and ever more complex IT solutions are being delivered. Extrapolating, we might wonder if the innovation cycle will become *disrupted* through IT security responses simply failing to keep-pace. Hopefully, like Moore's Law, we're on a continuum and this scenario will not trouble us in the near future?

Adrian Fern



Actifio

Founded	2009 Octifio
Founder(s)	Ash Ashutosh
Investment Funding	\$207.5 in 6 rounds from 7 investors
Website	www.actifio.com
Sector	Copy data virtualisation
Key Points	 Actifio's Virtual Data Pipeline[™] technology decouples data from infrastructure Replaces silo data management with a simple approach enabling production data to be used wherever it is needed Actifio helps global enterprises and service providers better utilise their data to drive business outcomes

Overview

Every time software creates data, whether for testing, compliance, disaster recovery, archiving, analytics or any other enterprise function outside the operational and business systems stack, it makes "copy data". This data grows at a much faster rate compared with end product data, which most businesses focus on managing. Copy data requires dedicated storage and management; Actifio offers an alternative, cloud-based management application that virtualises copy data and enables users to recover data instantly from the same storage location (thus saving silo storage and management costs).

Copy data consists of multiple copies of the same file. This data could come in the form of multiple copies of volumes, backups, test/dev, online copies for disaster recovery etc. An IDC study found that about 75% of storage is consumed by copy data. Actifio conducts joint research with IDC to understand the copy data problem in depth. IDC estimated that businesses would spend roughly \$44 billion on coping with copy data in 2013.

Technology

Actifio was one of the first companies to enter the copy data management (CDM) market. It keeps one "golden" copy of the data and makes it available for all company requirements without needing to store a separate copy for each use.



The system dramatically reduces unnecessary duplication of application data while ensuring it is protected and easily accessible through its entire lifecycle. By virtualising data management and storage, Actifio replaces client's silo data protection and availability solutions with a single platform.

Actifio's Virtual Data Pipeline (VDP) creates a single "golden master" of production data and maintains changes to that master copy, encrypting the data at rest and inflight. This enables users to run a simple "single-pane-of-glass" window application to recover large data sets in minutes.



Copy Data Virtualization

Applications

Some core use cases are:

- enable Enterprise Hybrid Cloud (eases the management of data and application workloads across legacy data centre(s) and cloud-based infrastructure)
- build higher quality applications faster (adds business value by making new capabilities available sooner; reduces deployment surprises and Dev Ops conflicts)
- improve business resiliency and availability (improve internal SLA's while reducing cost, risk, and complexity)



Actifio's products are:

Resiliency Director

An automated disaster recovery tool which automates the disaster recovery process from testing to execution

• CDS

A storage appliance that can scale to up to two petabytes; available in six standard configurations including services such as data lifecycle management, deduplication and backup replication, failover testing and fall-back automation

Gateway

An appliance with no native storage capacity that virtualises capacity for various vendors into a single storage pool that can be scaled up to eight petabytes

Observations

A lively and entertaining presentation was delivered by Kat Keegan (a16z Programme Manager and Account Executive), which prompted a good discussion amongst the tour delegates. Kat took away some additional questions in what has become known as "The Gerald Folder":

- Containerisation vs Virtualisation how is Actifio different from containerisation?
- HP3 Par how does Actifio compare?
- As data is being captured and put through the VDP to create the "Golden Master", how can we be sure that 100% of the data is moved from production into the Golden Copy?
- As data is ingested and processed at the block level, the Golden Copy is always going to be a few blocks behind production because of latency is this the case? If so, please provide further explanation.

Virtualisation of data completes the journey which started with compute, through network and storage, thus enabling complete Enterprise infrastructure virtualisation. Actifio claims its products cut data management costs by up to 90% and typically replaces five to thirteen separate pieces of software. The commercial model for products depends on the volume of data handled with rates costed per-terabyte, which decreases as volume increases.

Actifio says its clients get about £10 in cost savings for every pound spent on its products. TechValidate determined that: "Actifio's clients usually achieve 90% savings calculated on a total-cost-of-ownership basis."



Airware

Founded	2011 FAirware
Founder(s)	Jonathan Downey
Investment Funding	\$70.42m in 7 rounds from 15 investors Most Recent Funding: \$30m Series C on March 31, 2016
Website	www.airware.com
Sector	Drones, Robotics, Aerospace, Enterprise Software
Key Points	 Commercial drones help enterprises take advantage of aerial data for any application Airware's aerial platform combines hardware, software and cloud services to operate the drones at scale

Overview

Airware provides software that enables commercial organisations to make more efficient, consistent use of drones. Their solution incorporates cloud storage, software and drone hardware management; however, they do not supply drones, only the software to manage them, collect and analyse data. Airware provides the hardware integration, control software, workflow and analytics.

The Airware solution delivers software automation and enforced workflow to significantly increase the levels of drone safety and adoption for commercial uses, including insurance, utilities, telecoms and engineering.

Technology

Airware provides software that enables organisations to make consistent and efficient use of drones. In essence, the software is an auto-pilot for drones that makes them fully autonomous for specific use-cases such as surveys and inspections.

Using a device such as a tablet, a user manages the drone, for example where it needs to go and what service is required. Airware then fully controls the drone, following the selected plan to capture the images required which can be analysed using the Airware software.

Applications

The key markets for Airware are insurance, oil and gas, telecoms and utilities. Airware helps to reduce the cost of survey and inspection activities in any area where height



makes a specific site difficult to inspect due to accessibility / risk. Airware reduces the risk associated with untrained individuals piloting drones, both from a safety perspective and ensuring that the correct imagery is captured at the right quality. The software also provides some analytical capabilities such as allowing users to compare images taken at different times and highlighting areas of immediate risk.

Capturing images using a drone has some clear benefits around reducing the need for people to access difficult / dangerous places, reducing the cost of aerial surveys and providing imagery that may not previously have been possible. The other benefits include:

- Improved design and engineering assurance
- Lower infrastructure maintenance costs
- Enhanced safety and compliance

Observations

The Airware solution focuses on data needs rather than hardware. This allows organisations to make use of drones easily where currently the variety of hardware available and uncertainty about how to use them can be a barrier.

The use of drones is in its early stages, particularly as there is a limited amount of regulation around how and where they can be used. Software such as Airware that addresses the control-based risks of using drones by avoiding controlled airspace, controlling the flight path and giving the operator a safe and quick way of grounding the drone will help to make this capability more accessible. Airware are also working with aviation authorities such as the FAA, NASA and European agencies to help define the regulation needed.

Organisations will need to understand how they can use the imagery provided and are likely to need new workflows to incorporate both the survey itself and its analysis. There is also likely to be a requirement to protect the data to ensure that images gathered cannot be tampered with.



Arundo

Founded	August 2014
Founder(s)	Tor Jakob Ramsøy
Investment Funding	\$2.7M in 1 round from 3 investors
Website	www.arundo.com
Sector	Energy, Machine Learning, Artificial Intelligence, Internet of Things
Key Points	 Machine learning for assets and heavy industries Offices in Houston, Oslo and Palo Alto Provides analytical and predictive solutions that raise asset utilisation and performance in industrial companies

Overview

Arundo Analytics is a global provider of analytical and predictive solutions that raise asset utilisation and performance in industrial companies. Arundo's DeepQ product integrates sensor and transactional data to drive deep insight into historical equipment and maintenance performance. Arundo's real time predictive product, LiveQ, helps clients reduce maintenance costs, avoid unexpected outages and ultimately, improve revenue assurance through multi-layered use of patented machine learning techniques.



Arundo's SocialQ is a data driven collaborative problem solving platform that supports and leverages knowledge and expertise built up within heavy industries over many years. Its capabilities include:



- Rapid, data-driven problem solving across teams and companies
- Annotation of data, equipment and analytics
- Annotations using free text and labelling
- Metadata can be used for data mining at a later time e.g. for finding the time of failure or understanding team behaviour
- Advanced sharing options (enabling efficient collaboration)

-	
	-
Q	

Technology

Using a top down data driven approach, Arundo discovers the interdependencies between disparate equipment assets. The weakest link can bring down the whole system so it is essential that it can be identified to know what it is, when it might occur and how it will fail. The Arundo system also 'learns' from experience and operator interaction, ensuring that knowledge and competence is retained.

Arundo products are designed with advanced and secure technology 'from the ground up', delivered and updated from the Cloud. Using proprietary technology, it gives organisations the 'power of many', improving predictability by enabling the sharing of insights while keeping data secure.

Applications

Arundo Analytics' proprietary technology for asset-intensive industries uses Big Data and machine learning techniques to analyse industrial assets' data, in order to predict likely failures and improve operations. Asset-intensive industries are plagued with cost challenges due to unplanned downtime and periodic maintenance operations. Modern industrial assets are increasingly instrumented with myriad sensors constantly generating massive amounts of data, yet industry's ability to make use of that data has not kept up. In response to the market need for turning data into value, Arundo has developed advanced analytics products which lets organisations leverage the data sources they already have to provide actionable intelligence.



Observations

Arundo represents a new horizon in decision support, leveraging previous investments from solutions such as Alarm/SCADA (Supervisory Control and Data Acquisition), equipment specific condition monitoring and cross-asset analytics. Arundo's proposition is not about replacing existing integrated systems, but getting them all to work together to enable customers to maximise and assert control of all their data across all assets, equipment, sensors and knowledge bases.

Arundo is a great example of the use of machine learning applied to a real world situation. By concentrating in a specific domain of use cases, Arundo are able to develop and tune their algorithms to ensure that customers in their sector can continue to get maximum performance and continuity from their industrial assets.



Black Duck

Founded	2002 BLACK DUCK
Founder(s)	Douglas Levin
Investment Funding	\$71m in 7 rounds, plus 4 acquisitions
Website	www.blackducksoftware.com
Sector	Cyber security, Open Source
Key Points	 Identify use of Open Source components and lineage in software Understand security vulnerabilities, operational risks and licence
	compliance relating to Open Source use

Overview

Presented by Jukka Alanen, VP of Business Development and Corporate Strategy, Black Duck offers a solution to identify use of Open Source components within software. It provides insights into licence compliance, known security vulnerabilities and operational risks from redundant or abandoned code.



Black Duck refers to an increasing prevalence of Open Source code in mission critical applications within the enterprise. The diagram shows an accelerating trajectory to Open Source code becoming ubiquitous within the software codebase. Hence, it is increasingly important to understand and mitigate operational risks.

Perhaps somewhat surprisingly, Black Duck states that many organisations are

unaware that their applications may contain Open Source components. They highlight the diverse origins of software stacks, not just from outsourced code, but from other sources such as third party or supply chain components and reused code.

There are some real significant risks faced with organisations which have not identified use of Open Source code. This is not just because of the obvious security vulnerabilities which can be identified with any traditional application scanner or vulnerability testing, but also with licence (compliance) and operational risks.



Technology

The Back Duck suite is formed from three main components:

- Hub managing security and vulnerabilities
- Protex managing Open Source compliance
- Code Center managing governance of Open Source code

The solution can be integrated with development technologies, such as Jenkins, with plug-ins to further automate scanning and identification.



Integration With Existing Development Solutions

Black Duck provides different deployment models, including on demand, SaaS and onpremise.

Observations

As an example of Open Source compliance risk, GPL3 licence terms stipulate that the software can be modified and distributed, but any of these software modifications must be made available under the GPL.

Cisco is one example of successful litigation as a result of breaching these terms. They settled a lawsuit served by the Free Software Foundation in relation to a breach due to acquiring a Linksys product.

Black Duck's solution overlaps with other vendors we saw on the tour, though none of the solutions would cover the mitigation of both Open Source and security vulnerability risks. This presents a case for integration between at least two products to avoid over complication within the development lifecycle.



CaféX

Founded	2013 Café
Founder(s)	David Jodoin, Rami Musallam
Investment Funding	\$31.54m in 2 rounds from 3 investors
	Most Recent Funding: \$21m Series B on March 30, 2015
Website	www.cafex.com
Sector	Real Time, Enterprise Software, Telecommunications
Key Points	 CaféX is a leader in Web Real Time Communications (WebRTC) innovation CaféX delivers seamless and secure collaboration across applications on browsers, tablets and smartphones Personalised, intelligent connections

Overview

CaféX offer a set of digital engagement technologies that aim to improve customer and workforce engagement through making much better use of mobile and online channels, to help individuals to work together.

CaféX offers an easy to implement, frictionless video / chat experience that allows organisations to quickly improve their existing workforce or customer engagement capabilities. The solution comprises in-app communications, a one-size-fits-all video solution, and a contact centre performance management capability.

The aim of this is to make video collaboration easy to implement with other channels, enabling any organisation to make use of video as a way to engage customers or to enable employees to collaborate more effectively.

Their subscription model is based on the number of users of each product.

They have five separate products:

• Chime

A browser based video conferencing tool that does not require 'plugins' and works with any browser

Supervisor Assist



A contact centre utility that makes it easier to monitor and agents' performance through listening in, remote view and live webchat between agents and peers / supervisors

• Live Assist

Allows contact centre agents to collaborate with customers using screen sharing, co-browsing or app-control. It can be one or two-way and integrates with other contact centre applications

Kickstart

Allows non-technical users to build collaboration services very easily

• WebRTC in-app communications

Provides integration of communications and collaboration tools within existing apps using a consistent SDK across all platforms

Technology

Using proprietary technology, CaféX allows video-conferencing to run natively in almost any web browser without the need for additional plug-ins. The Chime application enables all web browsers to run collaboration suites without RTC, which is an existing barrier as it is not supported in some common browsers. Chime integrates with existing enterprise collaboration tools such as Skype for Business, Cisco and Zoom and allows many more users to take part through their browser. This can reduce cost and bandwidth requirements (mainly though combining multiple ports into a single conference session as it allows users on the same LAN to be clustered). The only requirement for the Chime tool to work is the availability of websockets.

Kickstart is a development tool that features a number of templates that allow nontechnical users to create and the compile code that embeds voice, video and cobrowsing in websites or apps.

The Live Assist technology is similarly independent of existing tools that can fail due to browser inconsistencies or lack of native app support.

Applications

The CaféX tool set has a number of use cases, the prime one being where an organisation needs easier to set up collaboration sessions, particularly in areas where desktop services are locked down so it is difficult to install new components or add-ins.

One particular case study for the collaboration tool set is driving collaboration between insurance company assessors and body-shops, using video collaboration to



complete an assessment and authorisation of repairs meaning that the overall repair process can be reduced by up to a week.

Observations

CaféX has a broad set of services, all of which are aimed at making the use of video easier. They have focussed on making video easy to implement and adopt by taking away much of the technical complexity of implementation and making it as close to plug-and-play as possible.

There is a clear requirement for tools such as this in the market as customers expect greater interaction, support and personalisation of services and employees expect and need to be able to work closely with each other independent of geography. Organisations need to be able to interact with customers seamlessly across multiple channels to ensure that the messaging, offerings and approach are consistent. CaféX fulfils some of these needs but is not a single point solution. Organisations must be careful to ensure they understand their customers and have a good understanding of the way customers want to interact before deploying point solutions such as CaféX.

The contact centre use case for the supervisor assist is more compelling given the increased need for expertise and high quality service within contact centres. Enabling managers to manage performance and interact live with both agents and customers remotely is likely to have clear benefits around service quality and retention.

The capability of the Chime tool to provide web-conferencing without any plugins and to integrate with existing solutions is interesting. Many organisations have already invested in tools such as Skype for Business, WebEx or Zoom, most of which need additional installations to work. The success of Chime is likely to depend on organisational attitudes to using apps versus browser based tools and on the actual impact of video-conferencing traffic on corporate networks. The traffic reduction offered by Chime may be the key selling point for this particular tool.



Cazena

Founded	2014	C A Z E N A big data as a service
Founder(s)	Prat Moghe	
	Vision: Big Data on demand	
Investment Funding	\$28m in 2 rounds from 3 investors	
	Most Recent Funding \$20m, Series B, July	2015
Website	www.cazena.com	
Sector	Big Data, Analytics, Cloud, BdaaS (Big Data	as a Service)
Key Points	First presented to Greenside delegates	during European Research
	Tour 2015, shortly after coming out of	"stealth mode"
	Late inclusion in the 2016 US Research	Tour due to a company no-
	show at Andreessen Horowitz	
	Cazena are an Ntegra Partner	

Overview

Big Data, combined with Cloud, promises great agility but, in reality, is hard to realise. Solutions are complex, not easily integrated (with existing systems), and rarely adhere to expertise class standards for security, identity management and compliance. Unpredictable costs and provisioning challenges also contribute to this complexity which, in turn, creates higher and higher barriers to entry. To enable these barriers to be lowered, Cazena's founding vision is to provide Big Data capabilities, on-demand, through their Big Data as a Service (BDaaS) orchestration and automation in the Cloud.

Technology

Cazena uses Infrastructure as a Service (IaaS) capabilities from a number of different cloud service providers, including Amazon Web Services and Microsoft Azure, to configure and orchestrate BDaaS environments.





These environments are provided in 2 categories: "Data Lake as a Service" for unstructured data and "Data Mart as a Service" for structured data. The solution provides a "Gateway Server", delivered as a Virtual Machine (VM) which the enterprise installs on their own infrastructure to allow data to be securely passed between existing on-premises systems and the Cazena service in the Cloud.

Web Console

Data Store		Care lever	Data Stilre	1	Cazeria Service
		-	Transfer Das Verticas +		Fight Data Mart 4
Contraction of the second seco	->	Carla More		•	Datest (Trining @rew
The Local File System					elstree
-0			1		×
Lance .		Contrast	lave		Contrast

Everything is managed through an intuitive and easy to use web console which enables users to offload and move data between environments with a few mouse clicks.

If required, enterprises can choose to have multiple gateways to allow collaboration or different access across various parts of a business (or with third parties). It is also possible to choose different geographical data zones to ensure compliance with data sovereignty rules.

Applications

The Cazena service can be used to store and manage entire organisational datasets or for more specific requirements. One use case is for organisations to use Cazena to store data about their customers that has been collected from external sources. These are stored in Cazena and can be combined with existing data held by the organisation



when needed to provide more insightful analytics. It is possible to ring-fence any data, in this fashion, with Cazena which also has the benefit of being able to tear down that data source if it is no longer needed without the usual decommissioning difficulties. A further use case is for Cazena to be used as a sand-box environment for new datasets or analytics. These can be quickly created and configured to explore different Business Intelligence (BI) approaches in a fraction of the time and effort that would usually be needed.

Observations

Cazena provides and orchestrates a single cloud service with multiple analytic processing engines, data movers, security and all functions needed for end-to-end processing. It is designed for fast, easy, secure integration with existing tools and on-premises systems through simple connectivity with the Cazena Gateway. As solutions and enterprise workloads evolve Cazena will benchmark, test and leverage best of breed technologies to transparently optimise and future-proof customers' platforms.



Crossdeck

Founded	2011 CROSSDECK
Founder(s)	John Moses, Joshua Awad, Josh Welle
Investment Funding	\$300k into a \$1.8m seed round
Website	www.crossdeck.us
Sector	Military, Mobile
Key Points	Workforce readiness reimagined
	 Daily work applications on secure platform
	 Digitise and mobilise military manual workflow and productivity

Overview

Crossdeck are currently targeted at the military with plans to expand to other government agencies and then to the private sector in the form of heavy industry e.g. oil and gas, and healthcare. The problem Crossdeck addresses is that 90% of the US military's workload is

As an officer in the Navy, responsible for hundreds of millions of dollars worth of equipment and the lives of sailors, this is the tool he was given to do his job

maintenance, repairs, training and safety operations. However, there is limited technology in use to help personnel manage and control their workloads.

Crossdeck's proposition is simple. They provide enterprise, daily work, collaborative applications for military customers on a secure platform.

Technology

Crossdeck has been developed using Microsoft Azure cloud hosted resources (end-toend) and deploys to standard Microsoft server and end-point infrastructure.

Applications

The current offering includes a suite of interconnected applications, including:

- Collaborative worklist
- Interactive "plan of the day"
- Mobile Personnel Qualification System



The roadmap includes Crossdeck as a platform acting as the access point for third party technology providers:



This software suite has applications for anywhere where mobility and security present difficulties for implementing enterprise collaboration tools such as document sharing, workflow and workforce management.

Observations

Crossdeck offers both a platform and a suite of applications. It has a clear business case based on the stated lack of digitisation in the US military. Further analysis may be needed to understand whether there are plans and contracts already in place to meet this need and whether this problem is reflected in other military forces or other private sector environments. It is also key to understand what capability is missing from the current large players (Salesforce, Microsoft, Box etc.), that drives the need for a new platform and set of services. If this gap is easy to fill, then Crossdeck may not be needed.

Similarly, integration with existing platforms and services will need to be considered and proven so that Crossdeck does not end up as a "dead-end" platform.



Databricks

Founded	²⁰¹³ \$ databricks
Founder(s)	Ion Stoicsa, Reynold Xin, Andy Konwinski, Patrick Wendell, Ali Ghodsi
Investment Funding	\$47M in 2 Rounds from 3 Investors (\$33M Series B in June 2014)
Website	www.databricks.com
Sector	Big Data simplification for the Enterprise
Key Points	 Founded by the creators of Apache Spark
	 Builds software and hosted services
	 Vision: to make big data simple for the enterprise

Overview

Databricks was founded by the team that created and continues to drive Apache Spark, a powerful open source data processing engine built for sophisticated analytics, ease of use, and speed. They are the largest contributor to the Open Source Spark project, providing ten times more code than any other company. To date, the company has the largest number of customers deploying Spark. Databricks is venturebacked by Andreessen Horowitz and New Enterprise Associates (NEA).

Technology

Apache Spark is an Open Source cluster computing framework originally developed at the University of California Berkeley Algorithms, Machines and People Lab (AMPLab). Spark provides an interface for programming entire clusters with implicit data parallelism and fault-tolerance. It was developed in response to limitations in the MapReduce cluster computing paradigm. Spark's Resilient Distributed Dataset (RDD) function is a working set for distributed programs that offers deliberately, a restricted form of distributed shared memory. The availability of RDD facilitates the implementation of both iterative algorithms, that visit their dataset multiple times in a loop, and interactive / exploratory data analysis, the repeated database-style querying of data. The latency of such applications compared to Apache Hadoop, a popular MapReduce implementation, may be reduced by several orders of magnitude. Among the class of iterative algorithms are the training algorithms for machine learning systems, which formed the initial impetus for developing Apache Spark.


Applications

Databricks provides an operationalised, enterprise ready, Spark platform that supports the transition from traditional predictive analytics / anomaly detection to truly data driven real-time analytics. It addresses the management of infrastructure and establishes Enterprise class "Production" configurations covering non-functional considerations such as High Availability, Disaster Recovery, Business Continuity and Enterprise Security with role based access to the environment integrated with Active Directory and / or LDAP.

The solution separates compute from storage, thereby enabling large volumes of data to be retained in, comparatively cheap, disparate cloud hosted storage to which compute resources can be spun-up and elastically applied on-demand. This "just-intime" data warehouse compute cluster approach significantly reduces TCO by ensuring expensive compute resource is only being paid for while it is being used.

Databricks empowers teams and improves their productivity through use of integrated workspaces "democratising access to Big Data", with customers reporting up to 50% productivity improvements.

As a guide, a PoC configuration and completion for 'one person, one use case' take approximately two weeks to implement.

Observations

A clear and concise presentation was delivered by John Tripier Senior Director of Business Development. This included a demonstration which suggested the orchestration of compute cluster resources to run against existing data repositories required no specific integration skills and could be undertaken by the same data analysts / data scientists that would be analysing the data.



Datameer

Founded	September 2009
Founder(s)	Stefan Groschupf
Investment Funding	\$76.75m in 5 rounds from 8 investors
Website	www.datameer.com
Sector	Data Integration, Data Visualisation, Analytics, Software, Big Data
Key Points	 Provides self-service, schema-on-read, Big Data analytics and visualisation of Apache Hadoop, for business users Scales up to thousands of Hadoop nodes Available for all Hadoop distributions Uses a 'spreadsheet' interface designed for business users

Overview

Datameer's end-to-end big data analytics application for Hadoop enables business users to discover insights in any data via wizard-based data integration, iterative pointand-click analytics, and drag-and-drop visualisations, regardless of the data type, size, or source. Datameer specialises in the analysis of very large (petabyte scale) data volumes for business users.

Datameer Analytics Solution (DAS) includes data source integration, an analytics engine with a spreadsheet interface designed for business users and visualisation including reports, charts and dashboards. DAS is available for major Hadoop distributions including Apache, Cloudera, EMC Greenplum HD, IBM BigInsights, MapR, Yahoo and Amazon.

Datameer uses Hadoop for both storage and compute, and is a fully extensible solution that can integrate with existing data warehouse or business intelligence solutions. Datameer users can also gain instant insights via the Datameer App Market, which offers dozens of pre-built analytic applications across a variety of use cases and data types.

Technology

Unlike traditional BI platforms, Datameer turns conventional IT-led processes into ITenabled processes to empower data-driven users across the enterprise. With selfservice functionality, more people across the organisation can discover new insights



in days instead of months and operationalise them immediately, increasing the agility and responsiveness of the business. All of this is enabled in a highly governed environment with enterprise-grade security and permissions functionality, "ensuring democratisation without introducing chaos".

Datameer is architected to empower analysts to easily load raw data of any type from any source into Hadoop and then run analytics directly on it using flexible schema-onread technology. Data remains on Hadoop for analysis, which means that no details are lost during the data transformation process and users can analyse their entire data set at once, not just a part of it. Also, the architecture streamlines the analytics process and ensures exceptional performance so users can get to insights quickly. As data never leaves Hadoop, full advantage can be taken of all that Hadoop has to offer. The flexibility of the architecture means it can easily incorporate new offerings in the Hadoop ecosystem, such as Apache Spark.

In contrast, many other big data analytics solutions require that analysts apply a schema to their unstructured data before analysing it. This forces them to guess which data attributes will apply to a given use case usually before they fully understand the value of the data value and what is actually needed. Often, this restructured data also has to be copied to external systems to perform data analysis outside Hadoop. This means users lose access to all of their data as they have to pick and choose what data to analyse which limits potential insights. In addition, structuring data and copying it between systems increases time to insight. As analysis happens outside of Hadoop users cannot take advantage of everything Hadoop has to offer.

Applications

Datameer is an enterprise-grade modern BI platform for agile analytics that allows organisation to answer new questions and identify patterns to help increase revenue, decrease costs, and dramatically increase overall business efficiency.

Datameer provides a model or framework that supports the modern data journey. It overcomes the current problems of: multi-structured data; complex analytical & correlation functions; meeting demand (productivity); reusing assets and skills; and working around the skills gap.

Observations

Traditionally, there have been very few people who have access to, and can use, corporate databases. Analyst friendly, Datameer seeks to 'democratise data' and enable many people, who should, to be able to access and use corporate data, particularly business users.



DeltaID

Founded	2011	
	(Product launch in 2015)	DELIAID
Founder(s)	Salil Prabhakar	
Investment Funding	Seed and Round A funding rounds, t	otalling \$6M
Website	www.deltaid.com	
Sector	Biometric Security	
Key Points	 Iris recognition for identity enro 	lment and authentication
	Better security performance wit	hout compromising functionality
	 Supplier to OEM providers 	

Overview

Presented by CEO Salil Prabhakar, Delta ID made a good case for use of iris recognition as a better biometric than face and fingerprint.

Face recognition is a challenging technology, with changes in lighting, hair and other looks resulting in a variable user experience. Fingerprints are a better and more popular biometric, but they can still be problematic.





Good, Bad, Ugly: 2%, 20%, 80% error!

Fingerprint is affected by weather, age, work, usage errors



Same person before & after fishing

Using the iris as the biometric, results in a much better accuracy.

US Research Tour 2016



	Devel	opment	Ma	iy 2013		
	6.05	***** ()#	es friger.	wateried		
	5.05		ex fager, or fagers	tobelled undathelised	···	
(MML) an	105	-25	ne fagers 1965, single	labeled a cyclicament		·
Non R	105	-20	vel, daf	iya camera	_	-
Tuthe Reject	2.0%	-10	teen, duilt oor fregerin	iya camera x . lebelod	_	_
	1.05	-				_
	0.05	1.06-4		1.05.05	1.05.04	105-0
		1,00,00		False Acceptance	e Rate (FAR)	

IRIS compared to fingerprints:

- lower false accept errors
- lower false reject errors
- improving false accept doesn't 'worsen' false rejects

The problem associated with using iris recognition is that it is a small object to capture. It requires high resolution imagery and with the continuous movement of eyes and short focus depth of field making it difficult to capture a suitable image. Delta ID solved this using off-the-shelf hardware, performing the 'heavy lifting' in the software to reduce costs and make it easier for device manufacturers to implement the technology.

Technology

The Delta ID technology is embedded in devices including those from Fujitsu, Samsung and Microsoft and is due for release in Q3 2016 on the HP Elite X3 Phablet, which is aimed at the enterprise market.

The product requires a two-megapixel camera with an infra-red LED. The infra-red LED provides more light and hence additional detail to prevent simulation or spoofing by use of photograph or other imagery.

For privacy, Delta ID confirmed that biometric data never leaves the device, therefore enrolment has to be performed separately on each device. The data can be used by services such as Microsoft Passport to allow single sign-on to numerous applications on a single device.

Applications

Delta ID illustrated various examples on phone / tablet security, including phone unlock and purchase confirmation / authorisation. A brief demonstration also showed that the authorisation was performed very quickly in a number of lighting conditions. Delta ID focus on selling to the OEM as additional hardware is required and so currently this cannot be implemented with software alone. The other consideration when integrating with device hardware is the "who boots first" issue. Ideally the



authentication software needs to boot first in order to ensure there is nothing of interest in the device memory until the user is authenticated and this is unlikely to be achievable with a retrofit software package.

Observations

Delta ID generated a lot of interest and questions on the day. This is reflected in the interest from automotive and IoT organisations who are looking for either a more robust biometric authentication mechanism than fingerprints, or they need something that can be used with gloved hands or hands-free.

Delta ID have found white space in the consumer and enterprise mobile market. They are likely to gain traction quickly as users look for a more robust, less obvious authentication mechanism than fingerprint technology. Iris recognition is already established in the government and military.

They have 25+ patent applications and so are very well placed to become the market leader in this technology.



Exabeam

Founded	2013 🦄 exabeam		
Founder(s)	Sylvain Gil, Nir Polak, Domingo Mihovilovic, Sylvain Gil		
Investment Funding	\$35M in 2 Rounds from 4 Investors		
	Most Recent: \$25M Series B on September 29, 2015		
Website	www.exabeam.com		
Sector	Cyber Security, Security, Information Technology		
Key Points	 Big Data security analytics Changing the way subgrattacks are detected 		
	 Simplifying security operations 		

Overview

Exabeam is a Big Data security analytics company that provides enterprises with advanced user behaviour intelligence. This unlocks the potential of existing security information and event management (SIEM) enabling IT security teams to quickly detect and respond to cyberattacks and insider threats in real time. User behaviour intelligence technology fundamentally changes the way cyberattacks are detected and simplifies security operations by focusing on attacker behaviour, rather than malware and tools, shining a spotlight on suspicious network activity and reducing the noise of false positive security alerts. Exabeam gives IT security teams the ability to detect and track credential use, in all parts of the attack chain, and prioritise responses through:

- extraction and enrichment of log feeds
- user session tracking
- behaviour analysis and event risk scoring

Technology

The Exabeam behaviour-based security intelligence platform combines advanced data science with powerful security analytics to create 'Stateful' User Tracking capabilities. Stateful User Tracking is what enables Exabeam to proactively hunt for threats, follow a user's tracks wherever they go in your network, automatically identify risky behaviour, and present risky profiles to security analysts, saving them hours and days in sifting through thousands of log alerts and false positives.





Exabeam connects individual user events into activity sessions. 'Stateful' User Tracking automatically stitches together users' activities into a distinctive session data model as they use different account credentials, change devices, and appear under different IP addresses. The resulting detailed timeline tells a security story about each session. Exabeam immediately identifies anomalous and 'out-of-character' behaviours, enabling accurate threat detection and accelerated response.

Security systems that are designed to detect, prevent, or alert on certain events quickly overwhelm analysts, making it impossible for them to know "Who is the user for this alert, what has he/she done since coming to the office, what happened after the alert, and is all of this normal?" Exabeam's 'Stateful' User Tracking holds state as each user changes credentials, devices, and locations over the course of the day to prioritise and deliver truly risky user profiles and activity using advanced statistical analysis with baseline profiling for deviation measurement. Analysis is based on categorical data, numerical data, and contextual information. Categorical data includes events that fall into specific quantifiable categories, such as the number of logons for a user from a specific country. Numerical data, such as number of assets accessed, duration of a user session, and time of day, is processed using real-time unsupervised clustering for discretisation. Contextual information provides additional insight, such as whether an asset is a workstation or server; whether an account is a human or service account; or if a device belongs to a privileged user. Context is estimated by multiple machine learning methods and helps calibrate and sharpen alerts.

Exabeam's techniques also support broader monitoring, such as cloud access, filelevel access, database table access, and application log monitoring. As data science



and security threats evolve, the Exabeam platform architecture supports new data science techniques to meet new security challenges.

Applications

Threat Hunter is an Exabeam security intelligence tool that uses 'Stateful' User Tracking session data models to complement user behaviour analysis. Exabeam Threat Hunter enables security analysts to search and pivot across multiple dimensions of user activity to find sessions that contain specific unusual behaviours or find users that match certain criteria. For example, an analyst might ask to see "all sessions where a user logged into the VPN from a foreign country for the first time, then accessed a new server for the first time, after which *FireEye* created a malware alert." This level of analysis across disjoint activities and systems is simple with Exabeam. Now analysts can ask new questions. With Threat Hunter, machine learning provides intelligent answers, in addition to alerts.

Observations

Exabeam provided an engaging and interesting demo. They're adamant that "simplicity makes for better security", and hide data analytics complexity for endusers' benefit. Although the rules exposed to users look simple, it takes exceptional design and engineering effort to give users the power to easily leverage deep, embedded security domain expertise and modelling sophistication that are critical for any data science-based effort. The result is data science output that users can easily understand and interpret.



GigaSpaces

Founded	2000	
Founder(s)	Nati Shalom, Eli Cohen	
Investment Funding	\$31m in 3 rounds from 5 investors	
Website	www.gigaspaces.com	
Sector	Cloud Computing, Software, Big Data, Virtualisation	
Key Points	 Provides software technology that enables organisations to automatically Deploy, Manage and Scale data intensive applications on any cloud. Brings data, applications and operations together Brings patterns from web-scale world to the enterprise 	

Overview

GigaSpaces Technologies, Inc. provides application platforms for Java and '.Net' environments. GigaSpaces has more than a decade of experience in developing software middleware for large scale real-time applications, bringing the best practices of web-scale architecture to enterprises through its two main products lines.

Cloudify is an Open Source Topology and Orchestration Specification for Cloud Applications (TOSCA) framework which brings the web-scale best practices of orchestration and automation into the world of hybrid cloud and Network Function Virtualisation (NFV).

InsightEdge and XAP bring web-scale best practices into the world of real time data processing and analytics.

The combination of the products is positioned to meet the demand of next generation applications in areas such as IoT and NFV in which insights from real-time data feeds and sensors turns into actions through orchestration.

Technology

Cloudify provides full lifecycle automation provisioning, deployment, monitoring and remediation of enterprise and NFV applications. It is a 'single pane of glass' across cloud platforms, containers and non-virtualised 'bare-metal', whilst supporting existing tools and install-base.



The Converged In-Memory Compute Platform by GigaSpaces combines extreme transaction processing with real-time analytics. XAP is a high performance In-Memory Data Grid that collocates data and business logic for ultra-low latency and high availability. InsightEdge is an In-Memory Analytics Platform leveraging Apache Spark for true real-time streaming analytics and actionable insights at web-scale.



Applications

- Ideal for complex, large-scale business-critical applications
- Orchestrates and manages any technology stack by using an open architecture
- Integrates any tool in the DevOps tool chain
- Native to OpenStack
- Enables OpenStack-VMware interoperability
- TOSCA compliant
- Supports all major clouds while keeping the application isolated from the underlying laaS
- Intelligent (data driven) orchestration



Observations

At the centre of two of the most important growing trends in modern enterprise IT:



Barriers to entry are high. It requires expertise in several IP intensive domains such as multi-tier applications management, extreme transaction processing and distributed computing.

Ovum, Jan 2015 "The more organizations use cloud computing, **the more they need tools that help them automate and manage** these various environments. Many vendors claim to meet this requirement. **Few can do so effectively**. While many tools are written for either a specific cloud or a specific capability, **Cloudify supports multiple clouds and capabilities**. Its new design and capabilities **make it a strong shortlist contender**, especially for those organizations interested in OpenStack."



Gigster

Founded	2013 GIGSTER
Founder(s)	Debo Olaosebikan
	Roger Dickey
Investment Funding	\$12.5M in 3 rounds from 12 investors
	Most Recent Funding: \$10m Series A on December 7, 2015
Website	gigster.com
Sector	Internet, Applications, Mobile
Key Points	 Experts in every area of the development stack
	Built web applications for Google
	 160m users for on-line products they have built

Overview

Gigster provides software development as a service. They provide access to a large, vetted and qualified community of developers who are able to rapidly and effectively create new software to a specific brief making good use of both human intelligence and AI to reduce timescales and improve quality. They call this a "tech team in your back pocket". This service is in response to the current development market which makes it difficult for organisations to hire developers and leads to problems with budgets, timescales and project abandonment.

Gigster aims to provide only top tier developers, those that have an excellent background combined with relevant experience in their domain. They also have a concept of karma – they look for developers to fit within the ethos of Gigster and to work well together as a team to ensure that the service provided always meets customer expectations.

Gigster functions in a similar way to most other development services in that they work from a specification, provide a quote and develop offsite with focus on regular communication and engagement with their client. They can build software either from a specification or from an innovation / problem statement provided by the client.

Gigster also provides ongoing support and maintenance contracts (as required) for any of the software that they have built.



Technology

Gigster functions as a service rather than a technology offering. They have a 'sweet spot' of areas where they have the most development experience, however this is driven mainly by demand rather than a desire to focus in a specific area so will change as demand changes.

They have an internal workflow product that helps to harness a geographically disperse community and makes use of AI / machine learning to understand how each developer works and assign projects / timescales accordingly.

Applications

Gigster have developed software for many global brands such as Mastercard, BCG, World Bank and Prudential as well as many other smaller clients. They most commonly develop applications in the mobile, IoT, data, machine learning an integration spaces.

Observations

Gigster is a service that is responding to a clear need in the development market, where developers increasingly want to work on a variety of projects for a variety of companies to increase their skills as well as keeping their interest level high. This makes it hard for organisations to hire full time developers which leaves them open to the risk of hiring on the contract market.

Gigster helps to reduce this risk by offering organisations the opportunity to contract with an organisation rather than an individual and also offers credibility that is higher than any one contractor.

The Gigster model or hiring the top tier developers and ensuring that they feel part of an organisational community because of their karma concept and because of the type of clients and work that they attract is likely to mean that developers stay with them for longer and therefore that organisations benefit from an increasing level of industry and individual client knowledge.

The Gigster model is likely to become the single most common method of hiring developers over the next few years and hence Gigster and its competitors will have to focus on ways to differentiate themselves to organisations looking to purchase as well as to individual developers.



Gigya

Founded	2006 GIGYA		
Founder(s)	Eyal Magen, Rooly Eliezerov, Eran Kutner		
Investment Funding	\$105.8m in 7 rounds from 11 investors		
	Most Recent Funding: \$35m Series F on November 4, 2014		
Website	www.gigya.com		
Sector	Identity Management, Computer, Enterprise Software		
Key Points	 Identity Management system offering flexibility and capabilities to deal with enterprise-grade security and data at scale Analytics, Auditing and Customer insights 		

Overview

The Gigya product is a Customer Identity Management platform. Customers now have much higher expectations around how easy it will be to interact with an organisation across multiple channels and around the level of knowledge the organisation has about them. The Gigya platform helps to manage and meet these expectations by creating customer intimacy and trust at the same time as providing the appropriate levels of security and governance.

The key premise that the Gigya platform is based on is that customers will give their personal information in exchange for personalisation, targeted offers or other perks. However, the existing way of registering and managing customers and the information provided is fragmented and the identity management solutions are often scaled for enterprises with up to 500k users rather than a customer base of millions.

Gigya supports organisations at three key stages:

- registration (registers and identifies customers across all sites and applications)
- profile management (aggregates and manages customer data which may be old or of low quality. Ensures that different regulations and consent models are delivered across different geographies)
- customer insights (integrates customer data into marketing and service applications to change unknown customers into known and loyal ones)



Technology

Gigya is a proprietary platform that offers integration with a wide variety of customer and information tools such as CRM, marketing, analytics and recommendation engines. This single repository of customer identity information helps to ensure that customers are understood by the organisation. All of their product information is collated, which traditionally has been difficult to link, along with their social / log-in details. This single view of the customer is made available to all required systems.

Applications

The Gigya platform aims to collect, connect and convert customers:



Case studies include:

• Forbes

Unified log-in, integrated with social and encouraged content consumption. Led to a 100% increase in new user registrations and 60% increase in comment activity

New Belgium Brewing

Increased registrations and application downloads which drove awareness of new products. Led to 8x boost in user registrations and 5x growth in average daily log-ins

Nike

74% increase in registration conversions

Observations

Customer expectations are increasing dramatically in line with the services, targeting and customer knowledge demonstrated by industry leaders such as Amazon. Customers expect to be asked to register in order to transact and often even to interact) with an organisation on-line but also expect that to be frictionless and that there will be reward in return for providing personal details.

US Research Tour 2016



Retailers, service providers, B2B organisations and any organisation with an on-line presence need to be able to meet these expectations consistently across all contact points and channels.

Identity management that is capable of integrating different internal customer identities e.g. for different products, with different external identities (email, social media, IM) so that there is a single record is critical in providing consistent services and in accurately targeting marketing and offers. This is also critical in terms of meeting regulatory needs around managing customer data and marketing preferences across borders.

There is still plenty of space for this market to evolve and it is likely that the existing CRM brands will continue to evolve their services in this space through development and acquisition. Organisations should understand the capabilities provided by their existing toolset before seeking to integrate something new but ultimately must consider the implications of not meeting customer expectations.



GoodData

Founded	2007 GoodData*
Founder(s)	Roman Stanek
Investment Funding	\$101.2M in 9 rounds from 17 investors Most Recent Funding: \$25.7m Series E on September 10, 2014
Website	www.gooddata.com
Sector	Collaboration, Business Intelligence, SaaS, Analytics, Enterprise Software
Key Points	 Open Analytics Platform Allows companies to manage and analyse their data in a seamless, interactive environment
	 Supports data governance, security and self-service data discovery

Overview

GoodData is a platform that enables the creation and distribution of analytics to multiple different entities within a business, such as franchises, partners, suppliers or stores. Their aim is to use analytics to create tangible business outcomes such as:

- improving solutions and relationships with partners
- increasing retention with existing customers
- creating data-driven products to create new revenue streams

They deliver this using a mixture of mature analytics, benchmarking and data enrichment to create and distribute useful business intelligence such as the impact of the weather on product sales.

GoodData offer both products and services in the form of the analytics and distribution tools, and the business architecture services needed to ensure that the right data products, based on valid information, are created.

Technology

GoodData is built around two key tools: the data analytics service and the distribution service. These combine to create useful insights and deliver them to the people who can make use of them to create and deliver business value.

It is a cloud platform that meets security and data standards such as ISO27001 and the Health Insurance Portability and Accountability Act (HIPAA).



The tools are licensed by an administrator rather than by users and are supported by professional services team with expertise in analytics across specific verticals

Applications

The GoodData service has applications across almost all verticals, anywhere that would benefit from analytics and insights being provided across a wide network of locations, roles and organisations.

In retail, as an example, the uses include:

- understanding regional store performance with retail sales, shopper engagement / loyalty programmes and website activity
- improving and tracking store opportunities by comparing how they are performing against their peers and making department adjustments
- analysing trends by department, category and products for better customer engagement, loyalty and improved margins

Clients such as McDonalds and Firehouse Subs have used GoodData to provide sales and other data to franchisees, driving performance improvements and better customer service through making the targets visible and providing insights to customer behaviour.

In the US Healthcare Information and Management Systems have created a data product using GoodData that they sell to organisational customers, for example trends in hospital spending sold to vendors of medical equipment

Other organisations across retail, financial services, hospitality and media have monetised data in this way.

Observations

Big Data and Analytics have been discussed for many years in the IT and business space, with limited success in implementing systems that deliver real business benefits. The success of GoodData is most likely down to the closed system offered, the combination of a product which is easily implementable and configurable with professional services expertise aligned to industries that helps organisations to get value from the tool. It is possible that the GoodData platform is not as powerful as other analytics tools but the overall solutions is more usable as it delivers benefits quickly and those benefits are easily understood.

This packaged solution is likely to be one of the key ways in which analytics tools start to gain traction and deliver real insights that change the way organisations use data, and make these insights available to smaller organisations. It addresses the key issue, that is, organisations don't know what they don't know. The expertise provided by GoodData helps to answer this question.



Joyent

Founded	2004 N.B. Acquired by Samsung Electronics, 16th June 2016		
Founder(s)	Dean Allen, Jason Hoffman, David Young		
Investment Funding	\$126M in 7 rounds from 11 investors		
Website	www.joyent.com		
Sector	Web Hosting, Cloud Computing, Infrastructure, Applications, Open Source		
Key Points	 Cloud infrastructure and big data analytics company offering a solution that powers real-time web and mobile applications Delivers container-native infrastructure offering organisations high-performance, yet simple public and private cloud software solutions for today's real-time web and mobile applications. 		

Overview

Joyent Inc. is a software and services company based in San Francisco specialising in application virtualisation and cloud computing. They build infrastructure software that automates data centre management and provides developers with a true cloud experience.

Nearly every industry is being disrupted by nimble, on-line and mobile challengers which increases pressure on companies to innovate quickly with software. Successful organisations can reduce development time from 18 months to two weeks by enlisting these approaches:

• DevOps

Development and Operational teams and tools merged to enable continuous deployment of software

Microservices

Legacy, monolithic applications being broken apart into discrete and more agile components

• Containers (Docker)

The preferred choice for packaging and deploying microservices styled applications.

"By 2018, more than 50% of new workloads will be deployed into containers in at least one stage of the application lifecycle" - Virtual



Machines and Containers solve different problems – Gartner, December 2015

To date, containers have lacked the operational tooling that exists for virtual machines. This lack of tooling impacts adoption, particularly in the Cloud. Cloud computing and containers make the perfect combination to enable agility and this is where Joyent fits.



Technology

JoyentCloud, Joyent's hosting capability, competes with Amazon's EC2 and offers Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) for large enterprises.

Joyent uses and supports open source projects including Node.js, Illumos and SmartOS, which is its own distribution of Illumos. SmartOS includes Joyent's port of the KVM Hypervisor for abstracting software from hardware, DTrace for troubleshooting and systems monitoring and the ZFS file system to connect servers to storage systems. Joyent is a founding member of the Open Container Initiative (OCI) and the Cloud Native Computing Foundation (CNCF), and works with Kubernetes and Mesos.

Container-native, Triton DataCenter converges container orchestration and cloud orchestration into one simple to deploy and manage solution. Alternative Container as a Service solutions add complexity and cost by layering container orchestration on top of cloud orchestration. Triton virtualises the entire data centre as a single elastic host capable of running Docker containers, infrastructure containers and Hardware Virtual Machines (HVM).

US Research Tour 2016





Applications

Joyent has taken their software stack which evolved to run their hosting business and is now licensing this under the name Smart Data Center to large hardware companies.

Triton's container hypervisor provides full isolation per container in a multi-tenant environment. Container infrastructure runs with the same performance advantages of bare-metal including lower latency access to CPU, storage and network resources, at data centre scale.

Triton enables users to mix Docker containers with instances of container-native Linux and traditional hardware virtualised machines to extend the benefits of containerisation to legacy applications and stateful services.



Observations

Docker-native tools make deploying on Triton as easy as running Docker on a laptop. There is no requirement to install or configure special software, nor is there complex host management. Triton automatically schedules containers across available physical hosts in the data centre.

Bare-metal provides performance but no scalability. Virtual machines deliver scale (multi-tenancy), but performance suffers because of overhead and resource costs. Joyent's Triton Container Infrastructure delivers bare-metal performance and the scale of a virtual machine, giving the best of both worlds.



Joyent describe themselves as being "Open Source but Opinionated" and provided a great presentation into a deeply technical area that has the potential to be quite dry.



Lookout

Founded	2007 Solotoot
Founder(s)	Kevin Mahaffey, James Burgess, John Hering
Investment Funding	\$282.34M in 7 Rounds from 19 Investors
Website	www.lookout.com
Sector	Cyber Security, Security, Mobile
Key Points	 Mobile device security platform Cloud based technology which is driven by a global sensor
	network
	• Combines machine learning and contextual analytics to protect mobile devices

Overview

Lookout is a security platform developed, specifically for mobile, in response to the increasing demand and expectation that employees can use mobile devices (a recent Ponemon Institute survey suggested that 33% of employees use mobile devices exclusively). Lookout offers mobile security services to both personal and enterprise users, with the main market being enterprise. They offer threat protection, a mobile security intelligence service and an API to make it easier to integrate analysis and security measures into a corporate mobile application workflow.

The Lookout mobile device security sits alongside other Mobile Device Management (MDM) solutions rather than replacing them. MDM solutions allow organisations to control apps and data on devices as well as to enforce security policies and deal with incidents. Lookout provides an additional security layer to this through allowing apps present on any device to be vetted, to manage any risks detected and to pro-actively look for new threats.

Technology

The Lookout service works through using their database of application codes (over 25m applications) and intelligent analysis to both identify existing threats on devices and predict potential new threats that can then be detected and defended against as soon as they are released. This database includes both applications in the popular application stores as well as applications that are often "side-loaded" on to both Android and IOS devices.



Lookout uses a combination of static analysis, dynamic analysis, reputation analysis and binary similarity to assess risks, to identify individual applications and known threats.

Applications

Lookout is offers both a personal and enterprise set of security tools. The personal tools include threat prevention as well as basic device tracking, backup and theft alerts.

The enterprise set includes a much more detailed set of security, MDM and threat prevention applications. Also, Lookout is developing a service that vets application store applications against corporate security policies to ensure that applications installed by users on corporate devices meet the relevant corporate policies.

To date, Lookout have over 100m installations of consumer client. Between these and the enterprise tools, approximately 7% of the 70k applications vetted by Lookout per day have security concerns.

Observations

Lookout research has found that 67% of the Global 2000 organisations have reported security breaches as a result of mobile devices accessing sensitive information. This level of threat has huge implications for both corporate and personal security given that there is a steady growth in employees wanting to use either corporate or personally owned mobile devices to access data and corporate services.

Organisations need to tread a fine line between ensuring that their information and data is protected at the same time as enabling employees (or customers), to work in a way that suits their role as well as meets reasonable personal preferences. The security concerns over employees using devices that may be uncontrolled on networks that also may be uncontrolled increase daily as threats to Android, IOS, Wi-Fi networks and other key parts of the access architecture continue to be revealed. Lookout offers one component of the security strategy that all organisations need to deploy but must also be recognised as being part of the answer to the security question rather than everything that is needed.



MarkLogic

Founded	2003 MarkLogic
Founder(s)	Frank R. Caufield, Christopher Lindblad, Paul Pedersen
Investment Funding	Total Equity Funding \$173.23M in 6 Rounds from 7 Investors Most Recent Funding \$102M Series F on May 11, 2015
Website	www.marklogic.com
Sector	Search Engine, Web Development, Content, Enterprise Software
Key Points	 A new generation database technology capable of handling any data, at any volume, in any structure MarkLogic Server is a NoSQL Database that enables enterprises to turn multi-structured data valuable, actionable insights Organisations around the world rely on MarkLogic's enterprise-grade technology to get to better decisions faster

Overview



What Business Wants:

• A unified, actionable 360 degree view of data

THE REALITY -> Data in Silos

- Data is spread across disconnected databases
- Mergers & Acquisitions outpace the speed of data integration
- Data needs to be delivered in real time



٠	80% of time wasted by data scientists wrangling data
٠	60% cost of data warehouse projects is on ETL
٠	\$36 billion spending in 2015 creating relational data silos!
	•

MarkLogic addresses the reality of data silos, providing capabilities that make it:

US Research Tour 2016





EASY TO **GET DATA IN** Flexible Data Model

Data ingested as is (no ETL)

- Structured and unstructured data
- Data and metadata together
- Adapts to changing data and changing data structures



- Index once and query endlessly
- Real-time and lightning fast
- Query across JSON, XML, text, geospatial, and semantic triples in one database



- Reliable data and transactions (100% ACID compliant)
- Out-of-the-box automatic failover, replication, and backup/recovery
- Enterprise-grade security and Common Criteria certified

MarkLogic is a multi-model NoSQL database that is able to store, manage, and search JSON/XML documents and graph data Resource Description Framework (RDF) triples. An RDF triple contains three components: the subject; the predicate; and the object. The predicate is used to denote relationships between the subject and the object and enables data to be stored and added to the database without the need to model complex relationships, in advance, using a schema or plan. For example:

Subject	Predicate	Object
MarkLogic	presented at	USRT2016
Sequoia Capital	provides funding to	MarkLogic
USRT2016	was hosted by	Sequoia Capital
USRT2016	was facilitated by	Ntegra Greenside

RDF is commonly used as a model to support semantics and its simplicity and flexibility means that otherwise difficult to model data relationships can be easily identified and thus dramatically reduce delivery timescales. Many organisations rely on the flexibility and agility of MarkLogic in order to integrate massive amounts of data and build large scale web applications.

MarkLogic is designed for *enterprise-grade* data integration.

Technology

MarkLogic has evolved from its XML database



ø

SCALE-OUT COMMODITY HARDWARE Scalable and elastic without over-provisioning or over-spending



SEMANTICS Enhance search and define data relationships



Integrate search with location and

geospatial coordinates



REAL-TIME ALERTING Unlimited alerts on new data at time of ingestion

roots to also natively store JSON documents and RDF triples. In addition to having a flexible data model, MarkLogic uses a distributed, scale-out architecture that can handle hundreds of billions of documents and hundreds of Terabytes of data. Unlike other NoSQL databases, MarkLogic maintains Atomicity, Consistency, Isolation,

Durability (ACID) consistency for transactions and has focused on building enterprise features into every release. Enterprise strength features include a robust security model certified according to the Common Criteria Evaluation and Validation Scheme (CCEVS), and enterprise-grade high availability and disaster recovery. MarkLogic is designed to run on-premises or in the cloud on Amazon Web Services (AWS).

Applications

During the 2012 London Olympics, the BBC used MarkLogic to power its Olympic Data Services. The application needed to be built in 12 months, which would not have been achievable using an SQL database which would have required the design to completely model the data.

The Results

- Scalable 2.8 petabytes of data delivered on peak day
- Fast 25,000 transactions per second, 45 billion requests
- Stable No outages or downtime
- · Success Personalized, award-winning Olympic coverage

BBC

BBC – Dynamic Content Delivery at Scale



The BBC broke all traffic records during the 2-week 2012 London Olympics Games, serving 2.8 Petabytes on the peak day, including more than 100m video requests.

Observations

Relational databases with ETL sacrifice agility, timeliness and cost. All future data needs must be predictable and new SQL queries require database re-indexing, and silo database changes require ETL re-writes. These considerations have been addressed fully by MarkLogic by developing a database that Integrates better, faster, with less cost.

It has flexible deployment options on Linux or Windows, On-premises (virtual or physical), private cloud, hybrid cloud, public cloud (AWS or Azure) using any storage architecture, including Hadoop.



Menlo Security

Founded	2013 Security	
Founder(s)	Poornima DeBolle, Gautam Altekar, Amir Ben-Afraim	
Investment Funding	\$35.5M in 2 Rounds from 4 Investors	
	Most Recent Funding: \$25M Series B on June 8, 2015	
Website	www.menlosecurity.com	
Sector	Network Security, Cyber Security, Security	
Key Points	 Eliminates malware from Web and email Isolation approach secures interaction online without deploying and point approach 	
	endpoint software	

Overview

Menlo Security protects organisations against cyber-attacks from the Web, email and other critical threat vectors. Their Isolation Platform isolates content and eliminates malware in the cloud, protecting users without the need to deploy endpoint security software. User deployments include large enterprises, Fortune 500 companies and financial services institutions.

Technology

The Menlo Security Isolation Platform (MSIP) is available as a public cloud service or for deployment in the user's data centre, the MSIP allows organisations to deploy isolation technology at scale to open up more of the web to users while simultaneously eliminating the risk of infection by web-borne malware.

MSIP eliminates the possibility of malware reaching user devices via compromised or malicious web sites or documents. The user's web session and all active content (e.g. Java, Flash, etc.), whether good or bad, is fully executed and contained in the Isolation Platform. Only safe, malware-free rendering information is delivered to the user's endpoint browser. No active content, including any potential malware, leaves the platform, so malware has no path to reach an endpoint and legitimate content needn't be blocked in the interest of security. Administrators can open up more of the Internet to their users while simultaneously eliminating the risk of attacks.





The MSIP is deployed between a user's device (desktop, laptop, tablet or smartphone) and the Internet. User Web requests are 'proxied' via the MSIP, which accesses the Web on the user's behalf and executes the user's session completely. Only safe, malware-free rendering information is sent to the user's endpoint, eliminating the possibility of malware reaching the user's device. The Platform, which provides isolation for both clear-text (HTTP) and SSL-encrypted (HTTPS) Web content, is available as a public cloud service and can also be delivered as a virtual appliance for deployment in an organization's data centre.

Applications

100% safety via isolation

Ground-breaking solution stops the never-ending search for risky content

• Seamless end-user experience

Safely empowers the digital workforce with a native user experience

• Cloud simplicity and scale

Reduces security complexity and increases scale by eliminating end-point software and out-dated appliances.

Observations

- Ransomware Prevention By executing user sessions away from the endpoint and delivering only safe rendering information to devices, users are protected from ransomware.
- Safe Access to Uncategorised Web Sites Architectures including a Secure Web Gateway (SWG) are often faced with a difficult decision around allowing access to new or less popular websites that have not been classified by the SWG.



Isolating uncategorized web sites via the Menlo Security Platform enables users to safely access more of the web while reducing the risk of malware.

- Safe Viewing of Web Documents The MSIP can eliminate the risks from weaponised documents (.pdf, .doc, .xls, .ppt) by isolating them in the Platform. Administrators can optionally allow users to download safe PDF versions of rendered documents (with all active content removed) and can also allow download of original documents for designated users.
- Eliminate Java and Flash from Endpoints Potentially harmful content such as Java and Flash is executed within the Platform, delivering a high-fidelity experience to the user without delivering any active content that can infect the endpoint. Administrators can remove Java and Flash on users' browsers but still allow access to Java and Flash apps without the risk of malware.
- Safe Email and Anti-Phishing The MSIP isolates and eliminates malware from sites accessed by clicking on links in emails. Additionally, the Email Isolation Service blocks user inputs to unknown web sites and thus prevents users from revealing their personal information.



Mesosphere

Founded	March 2013
Founder(s)	Ben Hindman, Tobi Knaup, Florian Leibert, Benjamin Hindman
Investment Funding	\$122.25M in 4 rounds from 14 investors
Website	mesosphere.com
Sector	Cloud Infrastructure and Data Centre automation
Key Points	 Mesosphere's Datacentre Operating System (DC/OS) is based on Open Source Apache Mesos (formally known as Nexus) Verizon selected Mesosphere's DC/OS as their nationwide platform for data centre service orchestration in August 2015 Microsoft announced a commercial partnership with Mesosphere in September 2015 to build container scheduling and orchestration services for Microsoft Azure

Overview



Mesosphere abstracts the entire datacentre into a single pool of compute resources, simplifying orchestration of, and running, distributed systems at scale. Mesosphere's Datacentre Operating System (DC/OS) is based on the open-

source Apache Mesos distributed systems kernel providing a common substrate for cluster computing and fine-grained resource sharing in the data centre. Mesos is the foundation for a number of distributed system frameworks used by Twitter, Airbnb and eBay. In April 2015 it was announced that Apple's Siri is using its own Mesos framework called Jarvis.

Technology

Apache Mesos is an open-source cluster manager that was originally developed at the University of California Berkeley RAD Laboratory. Mesos provides a scheduling mechanism that "offers" resources to frameworks, applications that run on top of Mesos. Each framework determines which resources to accept and what applications to run on them. This arbitration method allows near-optimal data locality when



sharing a cluster of compute nodes amongst diverse frameworks. The Mesosphere DC/OS spans all of the servers in a physical or cloud-based datacentre and runs on top of any Linux distribution. DC/OS provides container automation and governance for both 'stateful' and state-less distributed workloads including Docker containers, Big Data and traditional applications, in a single framework.



Applications

DC/OS is an enterprise class hybrid-cloud operating system for elastic applications that ensures high datacentre utilisation and makes executing and managing applications in complex environments as easy as launching applications on a laptop. It addressed the complexity of hyper-scale operations and supports developer and data agility.

DC/OS has a powerful Command Line Interface (CLI) and Graphical User Interface (GUI) to enable data centre operation, and a package manager to run containers and big data workloads in production.

Mesosphere provide two "turn-key" solutions that sit on top of DC/OS:

Velocity

Enables organisations to deliver modern applications in containers with scalable and reliable continuous delivery for development and operational teams

Infinity

An integrated product enabling companies to turn ubiquitous data to insight and action such as IoT, Predictive Analytics, Anomaly Detection and Personalisation

Observations

For a quick start and introduction to datacentre orchestration DC/OS can be downloaded from dcos.io

US Research Tour 2016



Mesosphere's easy to install, run anywhere, orchestration supports 'stateful' and stateless implementations with unlimited scaling potential. With its broad flexibility and maturity delivering a simple, single pane of glass, data centre management capability that leveraging the density levels achievable through the use of containerised solutions, it is very compelling. However, there is still a great deal of legacy investment in the heterogeneous point-solutions that populate many traditional datacentres of today.



Microsoft

Tour Host	1065 La Avenida, Mountain View, CA 94043Microsoft
Background	Microsoft was founded by Paul Allen and Bill Gates on April 4, 1975, to develop and sell BASIC interpreters for Altair 8800. It rose to dominate the personal computer operating system market with MS- DOS in the mid-1980s, followed by Microsoft Windows. The company's 1986 initial public offering, and subsequent rise in its share price, created three billionaires and an estimated 12,000 millionaires among Microsoft employees. Since the 1990s, it has increasingly diversified from the operating system market and has made a number of corporate acquisitions. In May 2011, Microsoft acquired Skype Technologies for \$8.5 billion in its largest acquisition to date.
Presentations	 Cloud update based upon their Microsoft Azure offering Machine Learning and Artificial Intelligence (AI)

Microsoft Cloud Update - Azure

Overview

Patrick Jean described the Microsoft Azure Cloud strategy (there were no slides for this presentation). He started the discussion by asking what questions were key to the group and structured the rest of the session around those. Questions raised were:

- What differentiates the Azure services from Amazon Web Services (AWS)?
- How does Azure 'handle' Active Directory (AD) multi-domain / geography capabilities?
- What is the future of Surface devices?

Technology

There was general agreement that users are now familiar with the concept of cloud services and are now keen to make greater use of services from Microsoft, Amazon, Google or other niche providers.

One of the key themes around the technology of Azure was security. Many customers / potential customers still have concerns about the security of cloud environments and Microsoft are keen to differentiate themselves in this area.



Microsoft assert that their built in security (such as DOS prevention) is ahead of the main competitor offerings and that, due to the scales of cloud provision security has more time and resource allocated to it than in an in-house environment. For example, the Azure environment provides the capability to audit and review every action taken by administrators, including those using the Office365 environment. Microsoft asserts that it is vital for companies to check up on administrators and that the use of cloud means now these roles are now filled by Microsoft staff rather than in-house staff. Microsoft suggest that the level of admin security and control and their level of certification is a differentiator.

The discussion continued to talk about defence in depth and the need for organisations to assume they have been breached and hence must focus on detection and resolution as well as prevention. The detection and resolution should include regular exercises to test how quickly the organisation (people, processes and technology) can respond to and recover from a breach. Microsoft, in common with other large organisations employ a red teaming approach (a dedicated team to look for security issues) to help identify and test breach processes.

Applications

60% of the Fortune 500 companies are already using Azure in some form and a rough poll showed that 3 of the tour delegates organisations already rely on public cloud for significant business services and others had plans to implement additional services within the next 12 months. Examples of this included internal video / YouTube hosting and delivery, B2B sites and marketing / public access services.

Cloud is now enabling more organisations to move towards platform as a service (service fabric). Demand is moving away from hosting towards services that allow organisations to deploy stable, scalable applications. One of the ways in which Microsoft is supporting this is to embrace open source software. This can be supported on Azure. In parallel, Microsoft is also considering moving to open document formats for Office.

Observations

Some of the issues still experienced when making use of cloud services include:

- organisations having to migrate services to the cloud, which is often painful versus greenfield implementations. This is possibly a temporary state as cloud services mature but organisations need to drive change by starting to design services for the cloud from the start
- storage in the public cloud is often seen as less performant than on premise.
 Whilst this can be true due to network limitations, it is not always the case and should not be an immediate showstopper for cloud projects


 lack of skills and knowledge around the use of cloud services – Microsoft suggests that organisations need to get the networking and storage teams to understand how to use and exploit cloud services first. Organisations do need to understand that using cloud infrastructure is a different skillset from managing in-house infrastructure

Limited discussion was held on the Surface product line. It was agreed that customers value both Apple and Surface devices (the assumption being that Apple and the MacBook are the major competitors for Surface products). Generally, it was felt that Surface is leading the market (over Apple) in some functionality e.g. soft keyboards but still has significant ground to make. It is not totally clear how committed Microsoft is to this product line, especially given the poor sales of the Surface Book outside of the US.

The last point made was that there is a future for both private and public clouds, including niche players e.g. connection to server such as Azure Stack / Flexpod.

Machine Learning and AI

Overview

Kevin Tsai presented Machine Learning and AI. He discussed the mechanics of machine learning as well as the importance of algorithms and the developing algorithmic businesses.

Much of the discussion focussed around deep learning and neural networks, which are one of the key tools in machine learning and their cost and capabilities are improving rapidly, year on year. The main problem associated with cognition, deep learning and analysis is processing capacity, the scale for this is non-linear meaning that processing more data equates to rapidly increasing costs that rapidly tend towards the prohibitive in today's world of Big Data.

One way to help address this is to use algorithms to reduce this cost by looking for rough answers initially (using less compute resource) to identify areas where further analysis is likely to be beneficial that are then fed into deep analysis.

This is one of the reasons that a market being created around algorithm selling in similar way as app stores (buy vs build). Algorithms for first stage processing as well as for deeper learning are being sold, including in the Microsoft machine learning marketplace (DataMarket).

In parallel, compute resource is also increasing as smaller components become more powerful. Neural networks are made of many small processors and so the cost and power per processor is key. The processors initially used were CPUs, now Graphics



Processors (GPUs) can offer more power at lower cost, the next step will be Field Programmable Gate-Arrays (FPGAs).

Technology

Kevin talked about the mechanics of machine learning. This covered transforming the data into something that the machine can understand, creating models such as parametric models and different types of learning such as reinforcement and stochastic learning.

One of the keys to deep learning is having a suitable algorithm. Many cutting edge algorithms are being developed by research institutions and are often made available as open-source. One of the sources of these – ImageNet – also has an annual competition to demonstrate the rapidly increasing capability of neural network learning.

Applications

Deep learning is an increasingly powerful unsupervised learning tool, where uncategorised data and fed in and the machine has to look for patterns and links to work out commonalities. A good example of this is the Google X-lab experiment looking at YouTube videos. The tool began to recognise human faces and cats with high level of accuracy after being presented with 10 million videos. The tool had never been given a definition or picture of a cat.

Applications of this in the real world include AI looking at oncology x rays to identify patterns and help to predict cancer; looking at data provided by heart rate monitors / pumps to predict problems or cardiac issues; providing analysis and stock market predictions or classifying / predicting anything that there is data for.

Observations

The capability and marketplace for AI and machine learning is continually developing, use cases are being identified daily. Machine learning is still in the early stages and its capabilities are not easily available to anyone except large organisations or developers with a lot of spare time. However, this is likely to change and deep learning will start to be both more pervasive and more available. This will lead to many changes about how we use computing and how much we trust machines to make decisions for us. This also needs to be tempered with more detailed knowledge about how the learning process works and how we can manage and control it.



Oblong

Founded	2006 •••••••••••••••••••••••••••••••••••		
Founder(s)	John Underkoffler, Kwindla Kramer		
Investment Funding	\$53.73m in 5 rounds from 4 investors		
Website	www.oblong.com		
Sector	Product Design, Hardware, Software		
Key Points	A design-driven software and hardware company		
	Develops interfaces to control windows across multiple displays		
	Basis for the computers depicted in the film "Minority Report"		
	• Mission is to put a new interface on every computer in the world		
	 Developed Mezzanine: Visual Collaboration for the Enterprise 		

Overview

Oblong Industries was founded in 2006 by John Underkoffler who designed the visual computer interface used in the film "Minority Report".

That science-fiction technology has now become a reality in Mezzanine, a visual collaboration solution that links locations, teams, content and devices in an immersive, shared workspace.



Oblong is a design-driven software and hardware company whose mission is to "put a new user interface

on every computer in the world". The immersive Mezzanine workspace spans screens, walls and locations with a simultaneous display of dynamic, computer-generated views, windows and images.

Technology

Oblong's technology originated from research undertaken at the MIT Media Laboratories. The "g-speak" Spatial Operating Environment (SOE) is a general purpose application development and deployment platform designed from the ground up. It gives programmers access to the principles of spatial and gestural computing and "real-world-pixels". This platform supports and enables development of applications that are inherently multi-screen, multi-user and multi-device.





Oblong's Mezzanine product, built on the gspeak platform is a meeting and conference room system that offers a unique shared collaboration space. Multiple participants work across multiple screens, sharing data, live video and data applications, exercising collaborative control from every device they bring with them. Telepresence capabilities, collaborative white-boarding and

presentation design and delivery are integrated into the Mezzanine experience. Each Mezzanine room can join up to three other Mezzanine rooms or connect to standard telepresence and videoconference endpoints, to enable distance collaboration.

Applications

Oblong has used g-speak to help companies like Boeing, Saudi Aramco, and GE as well as government agencies and universities. Profoundly productive solutions are built and deployed for modelling and simulation, logistics and supply chain management, natural resource discovery and extraction, energy grid management and financial and scientific data navigation and visualisation.

Observations

Oblong Industries were named a "Cool Vendor" by Gartner in 2016 for transforming Human Machine Interface in the Workplace. However, although our study tour delegates agreed the technology demonstrated at Oblong's Los Altos offices was very cool and visually immersive, it lacked the ability for users to interact directly with the projected images. If an application being displayed requires input, the user needs to find the device and interact with it locally. Oblong say they are working on this aspect.



Okta

Founded	January 2009		
Founder(s)	Todd McKinnon, Frederic Kerrest		
Investment Funding	\$229.25M in 7 rounds from 17 investors		
Website	www.okta.com		
Sector	Security, Enterprise Software, Telecommunications		
Key Points	 A foundation for secure connections between people and technology Allows people to access applications on any device at any time Enforces strong security protection Integrates directly with an organisation's existing directories and identity systems as well as 4,000+ applications 		

Overview

Okta (formally known as SaaSure Inc.) is a foundation for secure connections between people and technology. An 'Okta' is a meteorological measurement of the amount of cloud cover. By harnessing the power of the cloud, Okta allows people to access applications on any device at any time, while still enforcing strong security protection. It integrates directly with an organisation's existing identity directories and management systems, as well as 4,000+ applications.

Okta runs on an integrated platform which means organisations can implement the service quickly, on a large scale, at low total cost.

Technology

Okta, has created an integrated, cloud-based platform that securely connects any person via any device, to the technologies they need to do their work. Okta's products manage identity and access for millions of people on millions of devices. The platform provides:

- universal identity directory
- single sign-on
- automatic authentication
- provisioning
- centralised control of people and groups



- adaptive multi-factor authentication
- policy framework
- mobility management
- reporting and policy engine solutions

Okta offers IT products focussed on IT and security professionals. At a very high level, these products simplify the way people can connect to enterprise technology, while increasing efficiency and keeping IT environments secure. The Okta platform allows product teams and developers to layer Okta's powerful identity services into their applications and portals. The platform makes it easy to authenticate, manage, and secure users.

Applications

Thousands of customers, including Adobe, Allergan, Chiquita, LinkedIn and Western Union use Okta to help their organisations work faster, boost revenue and stay secure.

Observations

Okta's IT products use identity information to grant people access to applications on any device at any time. Their platform securely connects companies to their customers and partners and thousands of organisations trust Okta to help them fulfil their missions as quickly and as seamlessly as possible. They are 2016 Gartner "Magic Quadrant Leaders" for Identity and Access Management as a Service (IDaaS).

Also Okta provides professional services, support, training, adoption, and self-services.



Orbital Insight

Founded	2013	
Founder(s)	James Crawford	
Investment Funding	\$28.7m in 3 rounds from 8 investors Most Recent Funding: \$20m Venture on June 27, 2016	
Website	www.orbitalinsight.com	
Sector	Humanitarian, Non Profit, Financial Services, Analytics, Geospatial, Software, Big Data	
Key Points	 Satellite image processing and data science Innovative and interesting technique of getting insight from imagery both at home and outside of your own jurisdiction 	

Overview

Orbital Insight combines the wealth of publically available satellite imagery with image processing, machine learning and analytics to provide unique information and insights. This can provide insight into movements of people and machines, installations, geological changes, construction, illumination and anything that can be seen from above. The satellite imagery including non-visible spectrum images, can be combined with images from drones or other aerial imagery to provide key insights. The image information can be provided from a combination of current and archive images. Current images are most often obtained through a revenue share with satellite companies. Only 2% of imaging capacity is currently sold at the retail price based on specific, time-bound needs meaning that many have a lot of spare capacity. Orbital Insight takes advantage of this by buying imagery on a non-time sensitive basis to provide the core input for very detailed global analysis and insights.

Technology

Orbital Insight provides insight based on imagery then analytics using a combination of AI, machine learning and machine-human collaboration.

The capability is built in an Amazon Web Services (AWS) cloud and makes use of a number of Elastic Compute (EC2) instances running both Computer Vision (CV) algorithms to process images and Data Science algorithms to process numbers and/or volumes. The combined output from these delivers usable insight. It also uses "deep learning" techniques (looking for anomalies using machine learning / neural networks) to identify *unknown* unknowns.



Applications

Insights often start out as a search for the patterns of life such as "What does normal look like?" This can be analysed to look for anomalies and predictive analytics can be used to accurately predict future states. Also, this can start out as a wide area search where it is not clear what is being looked at or what users are trying to monitor. This can then result in new insights based on following individuals' lines of interest.

Orbital Insight most often work by partnering with organisations who have ideas, theories or initial analysis that can be proved or enhanced by the analytics and insights from earth imagery. Examples of this include:

- car counting at Walmart car park (this can be used as a proxy for people going to the store and so gives both current figures and can predict footfall. Used by analysts, this information can also give some insight into individual store performance which, when rolled up can be compared to retail predictions to spot gaps and invest in stock accordingly)
- identification of oil storage facilities (looking for the distinctive shape of oil storage tanks can easily be automated. The results can be analysed by assessing the level in the tanks to provide an accurate estimate of the total storage capability and the actual amount stored. This can provide an accurate view of the strategic reserves)
- predicting the rate of economic growth (this can be done through identifying and assessing changing levels of construction, illumination, car usage, pollution and others)

Observations

Big Data and Predictive Analytics have been discussed and hyped for a number of years as part of the big four IT topics i.e. cloud, mobile, social and analytics. However, analytics is only as good as the data you can provide to analyse and so a new source of data is always going to provide the capability for new and interesting insights. The use of satellite imagery combined with powerful image processing and analytics tools has the capability to provide highly relevant insights on a global scale. Although the company has only been inexistence for three years they have already provided insights that beats other predictions around retail sales and crop yields, both key areas for investors.

The amount of highly detailed imagery that is available means that there are many more uses for this raw data that have not been considered yet. With time, the right questions will be asked which will provide the means to exploit the imagery in the same way as other Big Data sources.



Origami Logic

Founded			
Founder(s)	Ofer Shaked, Alon Amit, Opher Kahane		
Investment Funding	\$49.3m in 3 rounds from 6 investors		
	Most Recent Funding: \$25m Series C on September 29, 2015		
Website	www.origamilogic.com		
Sector	Analytics, Big Data, Enterprise Software		
Key Points	 Marketing Signal Measurement Platform Providing clear real-time insight dashboard and exporting capability 		

Overview

Origami Logic is a customer insight tool that helps marketing teams make sense of data arriving from multiple sources. One of the key challenges with customer data is that the volume of data coming from multiple sources means that there is often a significant time lag between the event / campaign being monitored and the production of meaningful information and insight on the customer activity generated.

Origami Logic have developed application that can collect and analyse feeds from social networks or cloud marketing platforms in real time to generate meaningful insight quickly. This insight can be viewed within the Origami dashboard or exported into enterprise systems and applications.

Technology

The Origami Logic system collects information from multiple sources, currently external to the organisation but with plans to include internal data. This data is then cleaned (70% of the workload), validated and optimised so that it can be provided to an analytics team for further processing or to provide immediate insight.

Their "Marketing Signal Measurement" (MSM) platform runs a number of different algorithms to harvest, refine and analyse raw data from social media feeds e.g. Facebook, Twitter, LinkedIn, as well as marketing systems (e.g. Sprinklr, Nielsen). This information is then provided within a dashboard (that can be queried) as well as to other systems via an API.

The system can classify CRM aggregate data, not down to individual customer level, and also brings in data tags such as sentiment analysis.



Applications

Origami Logic principally helps marketing and products teams to make sense of the huge amount of social and customer data that is available to understand and track the impact of marketing campaigns, sponsorship or other promotional activity. The speed with which this is provided can help the teams to make changes and shape activity to maximise the impact. Typically, the system tracks activity such as Instagram or Facebook spikes.

Examples of where this has been used successfully include the Visa sponsorship of the 2012 Olympics where hundreds of millions of comments were tracked across multiple channels on a daily basis, where previously insight had not been available until after the two-week event.

Observations

Origami Logic offers fills a key need in the continuing requirement for organisations to stay in-tune with their customers. Customer expectations around both personalisation of services and speed of response have increased hugely in the last five to ten years and organisations increasingly struggle to both capture all of the data available and analyse it to provide clear information / insights.

The Origami product meets a large part of this need but needs to be integrated into a wider customer engagement strategy to ensure that the information provided is well managed and effectively used. It must factor in the customer information already held by organisations, often across multiple systems, as well as being effectively connected to customer service systems and product development lifecycles to be truly effective.

Organisations should also ensure that they have complete clarity around where the data is being sourced from to manage the risk associated with poor or incomplete datasets.



PinDrop

Founded	2011 Spindrop	
Founder(s)	Dr. Mustaque Ahamad, Vijay Balasubramaniyan, Paul Judge	
Investment Funding	\$117m in 4 rounds from 14 investors Most Recent Funding: \$75m Series C on January 28, 2016	
Website	www.pindrop.com	
Sector	Audio, Fraud Detection, Security, Risk Management, Mobile, Customer Service	
Key Points	 Anti-Fraud and authentication solutions for enterprise call centres Within 10 seconds a risk scoring is generating; reduces volumes of fraudulent calls, increases throughput of call centre agents and increased customer satisfaction 	

Overview

Pindrop addresses the growing issue of fraudulent calls to contact centres. Currently, one in every 2,000 phone calls made to a contact centre is fraudulent and each of these calls costs the organisation on average, 0.57 US dollars.

Addressing this fraud requires organisations to put in security measures such as voice recognition and challenge response (security questions). These are not only ineffective but they also worsen the customer experience meaning that genuine customers often get frustrated or worse when trying to make legitimate use of contact centres. This is in contrast to the security protection put in place for both physical and on-line channels.

Pindrop uses their technology called 'Phoneprinting' to examine all aspects of the call to score the potential for fraud. Low scores can be passed straight to an automated system or agent to service the customer. Higher scores can be passed through for additional security measures or to specialist agents to manage appropriately. This reduction in visible security for most customers combined with focussing on fraud likelihood enables improved customer satisfaction as well as reducing losses.



Technology

The Phoneprint technology used by Pindrop examines 147 factors about the call to generate a likely fraud score (not including voice recognition). These factors include:

- geolocation
- call type (internet phone, direct dial, etc.)
- unique phone (i.e. has the caller been identified before?)
- (dial tone modulation frequency) DTMF characteristics
- other risk factors

The system needs 10 seconds of information to build the print and generate a score.

Applications

Pindrop Phoneprint system is applicable for any enterprise scale contact centres such as banks, insurance, telco, retailers and government.

In 2015 \$8.6bn was lost in phone fraud attacks in the US alone. This contact channel is being focussed on by fraudsters because it is one of the easiest target channels due to the relative ease of bypassing the existing security and the volume of calls (18.9bn calls per annum in the US). The key security measure is currently testing knowledge through asking the caller questions that only the account holder / genuine customer should know. However, some research suggests that up to 50% of these questions can easily by answered by a fraudster based on research such as looking for publically available information e.g. from social networks, or through social engineering.

Additionally, the security questions are often asked by the contact centre agents themselves. They are incentivised to processes calls quickly rather than securely and so it is in their interest to ensure any call gets through the security process as quickly, rather than as robustly as possible. Together these facts offer a very large opportunity to fraud gangs.

The Pindrop service offers a way to reduce the opportunity for fraud as well as improving the customer experience. One specific example of the volume of fraud is chargebacks for retailers. Online retailers have estimated that 5% of phone orders result in a chargeback rising to 50% for premium brands such as Apple.

A further benefit quoted for Pindrop is associated with reducing the time taken for customers to complete their calls with the contact centre. The current security question approach typically takes between 45 and 90 seconds versus the quoted 10 seconds for a Phoneprint to be generated. This small time saving multiplied by the call volumes of a typical contact centre can lead to significant savings in the millions.



Observations

There have been many discussions in the last 10-15 years about channel shift (consumers moving away from face-face and phone contact to internet, mobile, social, TV and gaming device contacts). This now also includes chat-bots, automated responses to instant message conversations that can deal with everything from simple queries to sales. The way that this has involved is that consumers have embraced all of the contact channels, with preferences being shown for different channels by different social or peer groups. This has meant that organisations have had to develop multiple contact channels rapidly, leading to some of the legacy channels losing focus. One of the manifests of this is the approach to security for the telephone contact channel which is the area where Pindrop offers their service. There is a clear need for this both from a fraud reduction point of view as well as reducing customer frustration. The market for this is likely to be stable for a significant number of years as there is still a clear preference for using the telephone channel shown by surveys such as the annual Ofcom communications market report.



Perlert

Founded	2008	
Founder(s)	Stephen Dodson	
Investment Funding	\$11.34M in 2 Rounds from 3 Investors	
Website	www.prelert.com	
Sector	Computer, Software, Information Technology	
Key Points	 Provides incident management tools that automatically isolate the causality of application interruptions in real-time Helps IT teams find and fix real-time problems faster, using machine learning algorithms 	

Overview

Prelert Inc. provides incident management tools that automatically isolate the causality of application interruptions in real-time. One solution offered by Prelert is for business service management. This isolates the root causes of abnormalities in the behaviour of applications, in real-time. The Incident Cause Analysis software employs pattern detection algorithms to analyse terabytes of incident data. The company provides a causality analysis verification process that takes existing data and automatically identifies episodes of causality, such as streaming event, trend, and usage service management telemetry, which lead to application errors. It offers solutions for market data systems, Web-based applications, internal network applications, and on-line banking system environments.

Prelert addressed the key issues with behavioural analytics, for example, too much data, too many people, and not enough insight, to develop a "point and click" analysis solution that fixes the problems with software. They successfully applied their capabilities to security use cases, and have since been known as a cyber-security company.



Technology

Every day, terabytes of enterprise machine data is generated. Time-stamped logs from network devices, security appliances, servers, endpoints, applications, users, and more are retained, but the information required to effectively manage and secure complex IT infrastructures is hidden in all this data. It is simply too impractical for humans to find the "signal in the noise" unassisted. IT professionals have to develop and maintain a set of rules, signatures, and thresholds to monitor this data in an attempt to identify problems or incidents. Frequently, this leads to analysts and operators staring at meaningless dashboards plagued with false alarms.

Prelert helps IT teams find and fix real problems much faster, using machine learning algorithms that get smarter every minute. They do not require operators to write rules, create thresholds, or anticipate every possible move a hacker might make. Instead, Prelert analyses enterprise data, finds anomalies, links them together and tells the story behind advanced threat activity and operations issues. The algorithms it uses are based on enterprise data to ensure accuracy. The only way to know what is "abnormal" is to know what is "normal" for an organisation.

Understanding the Challenge

In this example showing a set of log events the questions to be answered are:

- what is "normal" or "abnormal"?
- what has changed and what were the causes?

1.1. DR. AN CALLADIT BRODUNDERLENDER UND RECEIVERLEISE (M. 1997) AND
and the day In Section on A. Les Jacober 1990 of 2014 and 100 day way and analysis of a section resources and the day of a section of the day of a section of the day
100-04-00 (2010-04) 04-00-04-00 (00-04-04-04-04-04-04-04-04-04-04-04-04-0
Michael (E) (E) (E) (C) (C) (C) (C) (D) (E) (E) (E) (E) (E) (E) (E) (E) (E) (E
HIGH TO THE TRANSPORT OF THE DATA OF THE DATA OF THE DATA OF THE
AND ADDITION TO A DEPOSE OF A DEPOS
Berlin and State and St
INT- AP (2010) PL 45 (ALL) PL
BHE-BY COLDUCT DE LETTER AND AND THE ME AND THE DE LETTER AND THE LETTER AND THE DETERMENT OF ADDRESS AND ADDRESS A ADDRESS AND ADDRESS
300-01-08 22 Little 1 00 Little 20 ND ND 200 ND
385-64-66 [20, 10-10] Web Math. Lett. 2012 [200 VD. ACT. 2013 10199 VD. 101 VD. 101 VD. 101 VD. 2012.] + 4, http://dxi EMDIT. ees. Instituted, and Vest. 7, http://dxi Vest. 7
BHE-BAR (D.N.) IN THE ARD DR. (TR) BAR (DR) BAR (DR) BAR (DR) BAR (DR) BAR (DR) CONTRACTORNAL (DR) BAR (DR)

The resulting operational issues can be:

- problems detected too late
- resolution takes too long



Applications

5.1; 101; All GA 1.1.6221" MIG	B Anti-Martineerin - JR. B. Th. W. E-479-5 magin Vew	And and
DNR Turneling American Sciences and Sciences Sciences and Sciences Sciences and Sciences Sciences and Sciences Sciences and Sciences	Proprietary post tree editories from 12.17 (20.220	099 Parties Transment
0 1 1		
Statistics of the second		

- Early Detection of Incidents
 - Detects advanced threat activity such as data exfiltration and "command and control" communications in near real-time
 - \circ $\;$ Identifies IT security and operational problems before users report them
 - o Catches issues that directly affect business KPIs
- Faster Root Cause Discovery
 - o Finds the root cause of anomalies faster
 - Gets the full story behind cyber-threats, IT operations issues, and business disruptions with algorithms that learn minute-to-minute what is normal for the environment
 - \circ $\;$ Involves fewer people in triage and gets answers fast
- Reduced False Positives
 - Prelert's analytics run on log data, from a broad set of sources, so they are able to consider more context than monitoring tools that rely on a single source
 - \circ $\;$ This additional context helps to significantly reduce false positives



Prelert deployment options include anomaly detection for Splunk and from Q2 2016, behavioural analytics for Elastic and API.

Anomaly Detective API Engine

- REST API & open source connectors
- Integrates to any data source







Observations

Prelert addresses the challenges related to understanding and baselining "normal behaviours" and hence, to be able to detect the "abnormal".

This is a great example of how machine learning can be applied to real-world situations, modelling time series data using "clever maths" and complex analysis functions. Prelet highlights the equivalence of Big Monitoring \equiv Big Data.



Skysafe

Founded	2015		
Founder(s)	Grant Jordan, Scott Torborg		
Investment Funding	\$3m in 1 round from 8 investors		
	iviost Recent Funding: \$3m Seed on April 20, 2016		
Website	www.skysafe.io		
Sector	Drones, Security, Drone Management, Aerospace		
Key Points	 Drone detection and notification with ability to selectively control individual drones (i.e. rogue drones) Provides detailed event history and real-time notifications 		

Overview

SkySafe provide anti-drone systems to public safety organisations. Their technology enables relevant organisations to take control of any rogue drones that present a safety risk. These can then be flown away and landed as needed.

This is in response to the growing number of consumer drones on the market and in use. These are increasingly causing issues for industries such as aviation, blue-light services, prisons and events as well as the possibility of being used for terrorist purposes.

The aim of SkySafe is to safely deal with rogue drones by controlling them and landing them safely rather than totally disabling them and forcing an uncontrolled crash-landing.

Technology

The SkySafe technology takes over the control of the drone from the original user by hacking into the control system, locking out the original user and making the drone respond to the SkySafe commands. Different levels of control are possible depending on the level of encryption used by the drone control system, ranging from full control by SkySafe to interrupting control by the original user.

The system relies on a series of hardware 'nodes' that surround the area to be covered by the system. These nodes both transmit the signals to the drones and allow the position to be determined. Each individual drone can be tracked and identified



through the content of the control packets. Drones can be uniquely identified and placed so that the correct drone can be selected for control. The content of the packets is also stored as data so that it can be used for audit / evidence purposes later as needed.

Drones can also be whitelisted to confirm that they are appropriately authorised to be in the air.

The software allows users to allocate and prioritise target drones following detection and identification. This is done at high speed.

There is also the possibility of using the system to prevent any drone usage in a specified area, although the way that this would be managed safely is unclear.

Applications

The main application for SkySafe is around public safety in areas such as airports, public stadia and emergency services work. There are many reported instances of drones causing disruption to air traffic or emergency services work and this often means that services needed to be changed or suspended to manage the risk while the drone operator is identified or the drone departs. The can cause significant impact to costs, timescales and endanger the public. Additionally, drones present a risk at public events where they could cause harm to the public.

A further use case is around protecting the security of prisons. There are a number of cases of drones being used to smuggle contraband items such as mobile phones, drugs or weapons into prisons. Many of these are thought to be undetected as many of those known are because the drone crashed.

SkySafe can be used to monitor and control airspace around specific areas to identify and track any drone usage and then to take over and control any rogue / unauthorised drones.

Observations

Like any security solution, the control of drones will be a cat and mouse game. The SkySafe service depends on them being able to break into the drone security and control encryption. While the drone market still has a huge number of players and is relatively immature it is likely that limited effort is being spent on the security of the link between drone and controller. However, as the market develops this could mature quickly. This is likely to then become an issue dealt with partly by national or international regulation. SkySafe will need to keep developing their "ethical hacking" capability in line with forging relationships with government agencies and regulators to help develop the relevant standards and ensure that they remain the software of choice in this area.



They will also need to ensure that their node hardware remains capable of broadcasting the signals needed to intercept and control any kind of drone.

A question was also asked in the session about the legality of taking control of a drone without the permission of the owner / user. This is not currently clear and may need additional regulations or laws to clarify what is acceptable.



Splunk

Founded	2003 splunk>	
Founder(s)	Michael Baum, Erik Swan, Rob Das	
Investment Funding	Went Public on April 19 2012 / NASDAQ:SPLK	
Website	www.splunk.com	
Sector	SIEM, machine data, IT Operations	
Key Points	 Monitor, report and analyse real-time machine data Comprehensible and actionable data reporting for executives outside a company's IT department Ability to correlate multiple large data streams, allow historical trend analysis and determine where real correlations between events and actions occur Google announced (2016) that it's cloud platform will integrate with Splunk, providing further opportunities 	

Overview

Doug Merit CEO of Splunk initially provided us with some insight into the company, before handing over to Nick Butler.

Splunk was founded 13 years ago and went public four years ago. Splunk is a platform used to search, monitor and analyse machine generated (big) data. With origins in SEIM and IT Operations Splunk has branched out by extending their technology to include a use of Hadoop. Their Hunk add-on provides Splunk Analytics for Hadoop, supporting accessing, searching, and reporting on external data sets located in Hadoop from a Splunk interface.

Splunk will collect data from anywhere, search and analyse anything and gain realtime operational intelligence. Splunk make a big thing of being able to "ask anything".

Splunk's unique selling point (USP) is that it can analyse data from almost anything, switches, active directory, event logs, door swipe data, etc. Using additional applications, this data can be compared pictorially and historically to determine trends and actions, without having to spend hours comparing lines and timestamps in epoch format. All this means that almost anyone can perform the analysis at any time, allowing the organisation to make rapid informed decisions.





Splunk is an established company with circa 2,000 employees, revenue of \$668m, representing year-on-year growth of 48%. It has 80 of the Fortune 100 companies as customers.

Technology

Schema on read was very much the message of the day. Splunk's ability to handle unstructured data and apply a schema on read means it only requires a time stamp on injection to be able to make sense and use of the data. All the effort is applied at index and search.

Applications

Multiple use cases for the reporting and analytics that Splunk can provide were outlined, including:

- application delivery
- IT operations
- security compliance and fraud (SEIM)
- business analytics
- internet of things (IoT)

Typical end customers within enterprises include not just security teams, but operations and application owners too.

Some examples of reference clients are:

- John Lewis for maintaining web uptime
- Gatwick Airport for resource planning
- British Airways
- ITV
- Airbus



Observations

When faced with a challenge regarding strategy relating to Open Source competition such as Elastic, Splunk maintain they want to retain the position as the "go-to" platform. Typically, they would workshop with potential customers the choice between "build vs buy", and invariably make a case for the latter!

The technology is obviously well established and has become an (IT) household name. It seems the focus is on how to gain additional penetration and to defend their position. It was noted that a significant amount of the presentation was focused on where Splunk technology is found, rather than on technical content.



Springpath inc.

Founded	2012 Springpath	
Founder(s)	Krishna Yadappanavar and Mallik Mahalingam	
Investment Funding	\$34M in 1 Round from 4 Investors	
Website	www.springpathinc.com	
Sector	Enterprise Software, Information Technology, Virtualisation	
Key Points	 Hyper-convergence software that turns standard servers of choice into a single pool of compute and storage resources Software defined storage stack on brand-name servers Recently partnered with CISCO 	

Overview

Springpath claims to have launched the first software-defined data platform that addresses virtualised compute as well as emerging container environments, in a hyper-converged solution. Based on hardware agnostic log-structure objects (HALO) architecture, Springpath's Data Platform is built 'from the ground up' to deliver elastic scaling, simple autonomous management, and a complete 'enterprise strength' feature set. The platform runs on industry standard x86 servers and empowers enterprises of all sizes to achieve an Independent IT Infrastructure, in data centres that benefit from silo-free environments, enabling businesses to stay agile and meet market demands.

Technology

Springpath's hyper-convergence software turns standard servers into a single pool of compute and storage resources. This Data Platform eliminates the need for network storage and intuitively integrates into existing management tools to maximize operational efficiency. Using adaptive scaling capabilities, customers can grow compute, caching or capacity resources independently, depending on their changing business needs. Springpath proactively monitors infrastructure to ensure resilient, always-on availability. It uses data management and optimisation capabilities to enable customers to experience transformative levels of resource utilisation, accelerating the adoption of DevOps in their organisations with truly agile IT



infrastructure.



Applications

Hyper-convergence solution focused on delivering enterprise reliability and performance on brand name standard servers.

•	Moving at the speed of business	(Consumption of infrastructure adds time, cost and inefficiency)
•	Operationalising speed and simplicity	(Proliferation of solutions creating silos that complicates operations)
•	CFOs expect Cloud Economics on	(Expectation of pay-as-you-go

 CFOs expect Cloud Economics of premise (Expectation of pay-as-you-go economics with scaling on-demand)

Springpath enables customers to deploy private cloud with public-cloud economy exploiting hyper-converged software running on industry standard servers. It delivers a wide range of data management and optimisation services at scale without impacting performance and eliminates silos with a single management platform for all applications and use cases running in any hypervisor of container.

Observations

- Previously known as Storvisor while in stealth. Emerged as Springpath in February 2015 after first round (\$34M from 4 investors including Sequoia Capital)
- Delegates were interested in CISCO's decision to partner with Springpath and wondered why it had not chosen to acquire it. This may be the logical next step.
- Subscription based commercials may prove difficult to align with corporate finance preference to capitalise own infrastructure projects.
- Centralised upgrade (when appropriate/necessary) could cause concern for IT operations as "Big Bang" changes are considered risky
- Provides secure segmentation allowing workloads of all security levels to exist in the same hyper-converged environment.



SS8

Founded	¹⁹⁹⁴ SS8 °
Founder(s)	Henry Wong
Investment Funding	\$40.5m
Website	www.ss8.com
Sector	Network security
Key Points	 Analysis of historic machine data against the latest security signatures Detect historic breaches with new intelligence

Overview

SS8 describe themselves as a time machine for breach detection - to go back historically and find breaches from the past. The point they are making is that security signatures are constantly being developed and rolled out. Although traditional intrusion detection and prevention service solutions look at information on the network for breaches happening now, there is no searches of past activity to find out if any of these breaches are detectable with the latest signatures.

Rather than storing network packets, SS8 summarises packet data by regenerating High Definition Records (HDRs) which summarise and enrich the content of the network traffic. This is analysed against the latest security signatures from third party threat intelligence to look for breaches.

The network is the key to detecting breaches and SS8 provide can detect and analyse network information to identify the Suspect of Interest (SOI) and Device of Interest (DOI). With a background working for law enforcement and intelligence agency customers, hence the SOI relevance, they have built their workflows and have expanded their customer base into the private sector for communications service providers and enterprise customers.

Technology

The key to the SS8 solution is the creation and storage of the HDRs, and not packet recording. HDRs are a summary and enrichment of layer seven network traffic, user and device information, providing information about the date and time, user, device, application, attachment and size.



Lightweight sensors are deployed on Internet points of egress. The learning engine and automated discovery platform can be deployed on premise or in the Cloud.



Applications

SS8 is designed to sit behind preventative tools, given its purpose is to detect rather than prevent or remediate.

Traditional security products raise the alarm when known threats are detected in the future. SS8is different in that it can identify when these known threats were exploited in the past and alert now and in the future.

A current breach detection model



The SS8 model:



Costs are based on volume of traffic analysed and length of retrospection.



Observations

Whilst this technology could be used as an improved alternative to some intrusion detection system (IDS) solutions, it would not negate the benefits of prevention or blocking and therefore would not displace an intrusion prevention (IPS) solution. SS8 and IPS could provide a powerful but ultimately disparate product combination. A question to answer is "Do I really want another security product in my estate?"

There is no disclosed plan to include blocking technology as this would require completely different architecture. Therefore, SS8 may be an acquisition target from the likes of Cisco, who acquired Sourcefire in 2014, to add IDPS capability to their portfolio.



Stack Overflow

(the flagship subsidiary of the Stack Exchange network)

Founded	2008
Founder(s)	Jeff Atwood, Joel Spolsky
Investment Funding	Stack Exchange funding: \$68M in 4 rounds from 15 investors Most recent round: January 2015, \$40m Series D from 5 investors including, lead VC, Andreessen Horowitz
Website	www.stackoverflow.com www.stackexchange.com
Sector	Communities, Recruitment, Internet
Key Points	 Stack Overflow: is a question and answer website for professional and enthusiast programmers receives more than 26 million unique visitors every month helping developers find answers to programming questions now hosts a recruitment platform "Stack Overflow Careers" which helps users connect with employers around the world

Overview

Stack Exchange is the largest network of websites for developers. It includes the subsidiary website Stack Overflow, a favoured destination for programmers, coders and developers to find, ask, and answer questions. The company was built on the premise that serving and enriching the developer community at large would lead to 'a better, smarter Internet'. Since then, the Stack Exchange network has grown into one of the largest collection of websites in the world, with Stack Overflow serving more than 26 million professional and novice programmers every month. The broader Stack Exchange Network has expanded to cover topics including mathematics, home improvements, statistics, and English language usage.





The website serves as a platform for users to ask and answer questions. Through membership and active participation, users are able to vote questions

and answers 'up' or 'down', and edit questions and answers, similar to using a wiki. Users of Stack Overflow can earn reputation points (for example, a person is awarded 10 reputation points for receiving an 'up' vote on an answer they have given to a particular question).



Users can also receive badges for their contributions (the 'gamification' of traditional Q&A forums).

The ability to close questions is a good differentiator from traditional sites (such as Yahoo! Answers) and provides a way to prevent 'low quality' questions. However, the community believes that duplicate questions are not a problem as they drive extra traffic to the site by multiplying relevant keyword hits in search engines.

The site celebrated its 10,000,000th question in August 2015 and, based on the tags assigned to questions, the top most discussed topics on the site are: Java, JavaScript, C#, PHP, Android, jQuery, Python and HTML.

A careers section assists developers in finding their next role. For employers, Stack Overflow provides tools to brand their business, advertise their openings and source candidates from Stack Overflow's database of developers, who are open to being contacted.

Technology

Stack Overflow is written in C# using the ASP.NET MVC (Model-View-Controller) framework and Microsoft SQL Server for the database. The 'Dapper' object-relational mapper is used for data access.

Unregistered users have access to most of the site's functionality. Users who sign in using the OpenID service, can access more functionality such as establishing a profile and being able to earn reputation points and badges, to allow functionality like re-tagging questions or voting to close a question.

User-generated content is licensed under a 'Creative Commons Attribute-ShareAlike' licence.

Applications

Stack Overflow boast that all of the world's developers and all of the tools needed to reach them, are available through their website. Although this claim may be



debateable, a "critical mass" of information, questions and answers, means anyone 'Googling' for development related information is almost guaranteed to find a valuable search result from Stack Overflow, prominently displayed on the first returned search page.

For example: http://www.google.co.uk/?q=xor+function+in+c

Stack Overflow is planning to exploit this ubiquity and 'monetise' their site in a number of interesting ways, exploiting for example, Insights, talent, advertisement and enterprise solutions.

Observations

One of the application ideas that Stack Overflow presented during the tour, was the ability to deploy a private instance of their website and database "inside a corporate firewall". This would enable an enterprise to protect and retain its own IPR in a closed community version of Stack Overflow that would immediately be familiar to new employees and hence decrease their "time to contribution" learning curve. Although this seems an obvious advantage, especially in organisations where there may be a high level of developer staff turnover, it was difficult to see how this model could be made cost effective.

Stack Overflow has received praise for its method of handling questions and answers as they are self-regulated by the community. However, it has been criticised in a study that found its policies discouraged women from being actively involved: bvasiles.github.io/papers/socinfo12.pdf



Striim

Founded	2012 Sstriim
Founder(s)	Steve Wilkes, Sami Akbay, Alok Pareek, Ali Kutay The Striim Platform was created by WebAction, Inc.
Investment Funding	\$42M in 3 rounds from 7 investors
Website	www.striim.com
Sector	Computer, Analytics, Software
Key Points	 End-to-end streaming data integration and operational intelligence platform Provides integration and intelligence enabling companies to make data useful "the instant it's born"

Overview

The Striim platform is an end-to-end streaming data integration and operational intelligence solution enabling continuous query / processing and streaming analytics. Striim specialises in integration from a wide variety of data sources including:

- transaction / change data
- events
- log files
- application and internet of things (IoT) sensor data
- real-time correlation across multiple streams.

Users can add: structure, logic and rules to streaming data; define time windows for analysis; detect outliers; visualise events of interest; and trigger alerts and automated workflows – all within milliseconds.

Striim is the only nonintrusive, 'enterprise strength' offering that combines streaming integration and intelligence in a single platform. Streaming data can be enriched with reference and historical data for instant context, at-speed and at-scale.

Technology

- Comprehensive end-to-end platform for both real-time data integration and streaming analytics
- Non-intrusive streaming change data capture (CDC) for transactional data
- Multi-stream correlation in real-time



- Enrichment of streaming data with reference / historical data for continuous integration and immediate context
- Declarative approach using SQL-like language to build and iterate streaming analytics applications rapidly; the entire solution can be built using a SQL-like language
- Continuously updating visualisations of streaming data
- Enterprise-strength and enterprise-scale leveraging low-cost compute
- Distributed In-Memory Computing (IMC) Architecture



Striim's IMC architecture comprises:

- in-memory databases (IMBDs)
- in-memory data grids (IMDGs)
- high-performance message infrastructure
- event stream processing (including complex event processing or CEP)
- data visualisation

This powerful combination of technologies allows it to support both real-time data integration and 'streaming analytics' use-cases without customers needing to piece together a 'mish-mash' of open source or proprietary solutions.

All of the processing is performed through event stream processing using in-memory continuous queries that can be written in an 'SQL-like' language (such as TQL). Intuitive user interface (UI) and TQL abstraction, with no coding necessary, makes the platform consistent and easy to use. These queries can carry out the filtering, transformation, and aggregation necessary for real-time data integration and perform the complex event processing, correlation and predictions necessary for streaming analytics.



Applications

With streaming data integration deployed, it is easy to 'layer on' streaming analytics to data pipelines, correlating streaming data with reference or historical data to immediately understand customers or business performance. Adding structure, logic and rules to streaming data and defining time windows for analysis enables operational processes to be triggered. Detection of outliers, visualisation of events of interest and raising alerts and workflows can be initiated within milliseconds.

Specific areas of interest include:

- customer experience (tracking and responding to top customer behaviours)
- location analytics (understanding and reacting to geographic movements in real-time)
- service level agreement (SLA) monitoring (tracking, predicting and projecting metrics)
- log Intelligence (correlation in context, in real-time)
- internet of things (at-speed and at-scale)
- security and risk (continuous protection)

Observations

The need for real-time insights is not isolated to a handful of industrial use cases. Any business operation can benefit from gaining access to data 'the instant it's born' and acting on it right away. So any company or business operation can benefit from streaming data integration coupled with streaming operational Intelligence.

Streaming integration should be included as part of any enterprise data strategy. Realtime analytics platforms, enabled by in memory computing, can support use cases by providing capabilities as an integrated set of "out-of-the box" features, and help contain implementation, maintenance, administration and support costs. Striim's presentation, during the tour, shows that they have created a platform that delivers all of these dimensions in a powerful and usable solution.



Synapse.ai

Founded	2015 Synopse.oi
Founder(s)	Early stage start-up still in stealth mode
Investment Funding	Supported by start-up incubator StartX
Website	http://synapse.ai/
Sector	Machine learning, Voice recognition, Mobile Technologies
Key Points	Siri/Cortana for enterprise mobilityTrademark "Rainmakr.Al"

Overview

Synapse.ai is an early stage start-up concentrating on improving the user interface for mobile professionals. Currently they are focussed on using voice to allow people to interact with their corporate systems, such as CRM and opportunity pipelines, while mobile.

Technology

Synapse.ai is based on similar speech recognition technology to Siri, Alexa, Cortana and other digital assistants. It integrates with existing corporate systems and uses machine learning / artificial intelligence to anticipate what information will be needed to both speed up interaction and help provide everything needed by the user. It also builds up a profile of individual users over time to make this process more effective, thus predicting questions and requirements more effectively over time.

Applications

The primary application for this technology is being used to support mobile business development & sales employee. The Synapse tool is integrated with existing CRM and pipeline systems to provide easy access to information about the contact, the opportunity, previous conversations and other information from the corporate tools. This is all provided through an intelligent voice interface to enable the employee to converse with the tool rather than have to present a series of isolated questions.

Wherever users need to interact with corporate systems, in some specific environments, there are applications where screen-based UIs would not be ideal. Some examples are driving a vehicle, conducting maintenance work, making deliveries and many other areas where hands / eyes are already occupied.



Observations

The Synapse tool is at the seed stage and so it remains to be seen whether it will be successful, as a stand-alone tool, or is likely to be bought by an existing digital assistant or corporate tool provider to enhance their capability.

It is clear that the future of UI is not just about text and keyboard / touchscreen entry. Many analysts are predicting a jump in the use of voice search and voice interaction as well as other non-contact methods such as eye tracking.

It is likely that the next big development in technology will be around humanisation, making the way that we interact with technology more natural and less forced. Synapse Voice is one such tool and takes the existing capabilities of consumer tools both into the enterprise and into a greater level of interaction based on a specific use case. Developing a wider set of use cases in combination with intelligent learning will enable voice interaction to become much more pervasive.


Talena

Founded	2013 Talena
Founder(s)	Nitin Donde
Investment Funding	\$12M in 1 Round from 3 Investors
Website	www.talena-inc.com
Sector	Software, Big Data, Database
Key Points	Company currently operating in "stealth mode"
	 Next-generation data management platform aimed at disrupting big data and the storage industry.
	 Always on Big Data
Founder(s) Investment Funding Website Sector Key Points	Nitin Donde \$12M in 1 Round from 3 Investors www.talena-inc.com Software, Big Data, Database • Company currently operating in "stealth mode" • Next-generation data management platform aimed at disrupting big data and the storage industry • Always on Big Data

Overview

The Talena platform evolved to address today's big data management challenges and ensure 'always-on data'. The Talena platform optimises test and development management, backup, recovery, and archive functions. Users can rapidly iterate application development and usage, prevent data loss, and minimise compliance risks.



Organisations are increasingly faced with managing information streams of exceptional volume and complexity, and are in need of more powerful and efficient tools for capturing, storing, organising, securing and analysing data to gain business advantage, with 'big data' becoming critical.

"Business-critical data is behind every smart decision, every success, and our customers understand that," said Talena's Founder and CEO.



Technology

Talena launched its latest innovation, Talena ActiveRx[™] in March. A predictive analytics infrastructure, ActiveRx, incorporates machine learning into big data workloads to accurately predict data availability and turn idle backup assets into active data copies. The Talena cluster is a highly available, storage agnostic, distributed software product that supports internal drives, direct attached storage and cloud storage in hybrid aggregated models. It runs on industry standard x86 physical or virtual servers, on premise or in the cloud.



Manage Data Availability Across Application Lifecycle

Applications

Talena is purpose built for Big Data management, providing a "single pane of glass" user interface (UI) across multiple sources, for multiple use cases, achieving enterprise strength operational requirements, such as: RBAC; LDAP integration; throttling; and blackout windows. This enables organisations to deploy and manage complex, high value big data applications, such as: NoSQL; Hadoop; and Modern Enterprise Data Warehouses, while conforming to traditional operational disciplines & requirements.





Observations

- No scripting required
- Free up storage on primary systems
- Reduce backup storage consumption
- Retain more restore points
- Automatically delete older restore points
- Simpler and faster restores using Talena (FastFind, "virtual fulls")
- Easily restore to a cluster of different size
- Network efficient incremental forever
- Integrated backup & test/dev management with masking and sampling
- Commercial model is based on capacity pricing