From Big Data to the Big Picture

Orion Health White Paper
David Bennett
VP Product & Strategy
072016

ORION HEALTH

The power of an Open Platform in Healthcare

The healthcare landscape has rapidly changed in the past few years. Healthcare spend globally has been increasing as a proportion of GDP and the outlook is for further upward pressure due to ageing population and the increased incidence of chronic illness. Focus in care delivery is shifting from a fee for service model to proactively managing the health of a population to achieve the Triple Aim: improve health outcomes, improve healthcare quality and lower healthcare costs.[1] Fundamental to Population Health Management (PHM) is an information technology infrastructure for the acquisition and aggregation of health data to provide longitudinal patient records; information which can then be analysed to provide actionable insights. The foundation of this is a system to manage, store and process the volume and velocity of health data, from multiple and growing sources, including claims, clinical, consumer and social data. The healthcare system of the future requires a solid IT infrastructure today.

A healthcare system in crisis?

Healthcare spending globally, on a macro basis, has been increasing rapidly for the past few decades. In the US, healthcare as a percentage of GDP has doubled in the last thirty years.[2] The Centers for Medicare and Medicaid Services (CMS) estimate that US health spending was 17.2% of GDP in 2012 at US $2.8tn and will reach US $5.2tn, or 19.3% of GDP, by 2023.[3] Australia's health expenditure has increased from A$82.9 bn in 2001-02 (8.4% of GDP) to A$140.2 bn in 2011-12 (9.5% of GDP).[4] Many developed countries are now spending above 10% of GDP on healthcare.[5]

It's no surprise that healthcare costs are in crisis, globally. The drivers of increasing health spend are complex and multi-faceted and pervade almost all countries. First, the world's population is aging but living longer. Studies have shown that healthcare per capita spend accelerates from around 60 years of age[6] and the World Health Organisation projects the proportion of the world's population over 60 years of age to double from 11% in 2000 to 22% by 2050.[7] Advances in science and medicine can keep people alive longer – a desired, yet costly, outcome for healthcare systems globally.

Second, there is a significant increase in the reported incidence of chronic disease, which results in greater costs for a health system. Chronic diseases are characterised as long-lasting conditions that can be controlled but not cured, such as diabetes, stroke and heart disease. The World Health Organisation states that chronic diseases are by far the leading cause of mortality in the world, representing 60% of all deaths.[8] The Australian Institute of Health and Welfare estimates around 82% of people over 65 have one or more chronic disease.[9] In the US it is estimated that 5% of the population account for 50% of the healthcare budget, due to the prevalence of chronic conditions.[10] Nearly one-third of US adults are obese and nearly half of adults have a chronic illness, including heart disease and diabetes.[11]

---

[1] http://www.ihi.org/Engage/Initiatives/TripleAim/pages/default.aspx

[2] World Development Bank, data

[3] https://www.cms.gov

[4] Review of the Personally Controlled Electronic Health Record, December 2013

[5] World Development Bank, data

[6] http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1361028/

[7] World Health Organization

[8] http://www.who.int/chp/en/

[9] http://www.aihw.gov.au/chronic-diseases

[10] http://www.nationaljournal.com/healthcare/report-5-percent-of-people-account-for-half-of-u-s-health-care-spending-20110627

[11] http://www.hcup-us.ahrq.gov/reports/statbriefs/sb144.pdf

Third, hospital visits are incredibly expensive. The average costs per stay in the US increased 45% to $9,700 in 2010.[12] Certain hospitals around the world are now sending "well" inpatients, such as C-Section recipients, to five star hotels for their recovery because it is more cost effective than occupying a hospital bed. Reportedly, many hospitalisations are unnecessary. According to the latest News and Numbers about hospitalisations from the Agency for Healthcare Research and Quality (AHRQ), approximately 1 in 10 of the nearly 40 million hospital admissions in 2008 were potentially avoidable.[13] These admissions were primarily for conditions such as diabetes, dehydration, and certain heart conditions and infections, which can be treated with appropriate outpatient care.

Fourth, the lack of a coordinated healthcare system not only results in unnecessary hospital admissions but also increases the risk and probability of medical errors. These carry significant costs. The Inspector General of the US Department of Health and Human Services has estimated that 13.5% of hospitalised Medicare beneficiaries experience adverse events during their hospital stays, of which 44% were preventable.[14] Additional Medicare costs associated with the adverse temporary harm events cost Medicare an estimated 3.5% of its budget in 2009, equating to US$4.4bn.[15]

The crisis in healthcare, caused by the trend in burgeoning costs, needs urgent attention and rectification. Globally, the focus needs to be on improving the efficiency of healthcare spend while improving health outcomes and the patient experience. We need to keep people healthy and out of hospital. We need to focus on managing the health of a population, which requires having access to all of an individual's health information and a system for turning that information into action. We need an information technology infrastructure and a focus on greater IT adoption and integration.

Big data, big deal?

Population Health Management moves the focus of healthcare from being responsive – providing care at the point of need – to being holistic, proactive and patient-centric through efficient care coordination and patient engagement. It is recognised that a move to PHM requires a high degree of information sharing between care providers, patients and their caregivers. But critically, it requires extracting and aggregating all the information about an individual, to provide a comprehensive patient health record.

The data we need to understand the health of an individual, and population, comes from multiple places and is of varied type. First, clinical data is fundamentally important to gain a deep understanding of an individual's health. Traditionally this information has been captured in siloed information systems, with only minimal sharing between silos, and only minimal amounts of information shared that the sender believes the recipient needs. This results in a number of harmful effects; from sub-optimal care delivery if important information is not known (such as a medication allergy) to resource wastage such as the duplication of laboratory tests. Pulling all clinical data into one central repository to provide a single record of data about a patient, available to all who need to access it, will provide immediate benefits both in terms of the quality of care that can be delivered and cost improvements by eliminating duplications and avoiding unnecessary emergency department visits.

[12] http://www.hcup-us.ahrq.gov/reports/statbriefs/sb144.pdf

[13] (http://archive.ahrq.gov/news/newsroom/news-and-numbers/110310.html)

[14] http://oig.hhs.gov/oei/reports/oei-06-09-00090.pdf

[15] Ibid

But clinical data is just part of the picture. Claims data, from health insurance organisations, is important to combine with clinical data to provide depth of information about an individual. Claims data is administrative data and includes information about patient demographics, billable charges, dates of service, diagnosis codes, procedure codes, prescription medications, insurance and providers.[16] This data is generated after every patient encounter with a provider that is subsequently submitted to a health insurer (being a government or private insurance company) for reimbursement. Traditionally claims data has been useful for analysing mortality rates, complications, and charges for care provided, but on its own, claims data has limited use for quality and cost improvements.

But clinical and claims data is not enough. To keep people out of hospital, we need to understand their conditions and how they are tracking right now so that care providers can intervene early when signs of deterioration appear. This necessitates real-time (or close to real-time) monitoring of individuals' key vital statistics, such as glucose monitoring for diabetics, uric acid levels for gout sufferers, and blood pressure for a patient with major cardiac risk profile. And with the growth of popular wearables and the connected "Internet of Things" it is now becoming possible to capture this information on a large scale. Such scale, however, brings some challenges.The IT infrastructure must be able to deal with massive volumes of small data from millions of patients. And it is highly likely that device data will have significant highs and lows in terms of data submission. At night, when device input is low, you want the system to scale back and consume fewer resources, but you need to be able to scale up for times of peak activity – such as the last few minutes of a close Superbowl. Therefore, you need a system that is elastically scalable – one that can adapt to the load and automatically optimise itself.

There are additional types of data to consider also, including machine-to-machine data (such as telemetry data), big transaction data (such as financial transactions), human-generated data (including unstructured and semi-structured data from EMRs, physicians free text notes, and video recordings), as well as environmental data (about pollution, chemical exposure and emissions). Social data is also being recognised as a significant factor in predicting long term health – information regarding abuse, neglect and stress that can have a real, tangible effect on an individual's physical health and wellbeing. Plugging this data into an aggregated record provides a truly comprehensive longitudinal health record for an individual and the complete picture for any care provider, providing the ability to provide decision support at the point of care.

Combining all forms of data into a single longitudinal record – a single source of truth – is optimal. But it means that we need the ability to handle extremely large amounts of data. Data Volumes are already very high. Data from the U.S. healthcare system alone reached 150 exabytes in 2011.[17] At the current rate of growth, big data for US healthcare will soon reach the zettabyte scale (1021 gigabytes), followed by the yottabyte scale (1024 gigabytes).[18] Kaiser Permanente, the California-based health network that has more than 9 million members, is believed to have between 26.5 and 44 petabytes of potentially rich data from EHRs, including images and annotations.[19]

---

[16] https://www.healthcatalyst.com/white-paper/clinical-claims-data-population-health-management

[17] http://www.hissjournal.com/content/2/1/3

[18] Ibid

[19] Ibid

We are in the age of Big Data. And it is a big deal. Genome sequencing is still in its infancy but promises to become more accessible as it is normalised into healthcare practices. The first genome sequencing took about 13 years to carry out, but now the largest labs can sequence human genomes at the pace of two per hour.[20] The data volumes for genomic sequences are huge – in the order of 1GB per person and up to 1TB in some cases – so the volume of data that must be analysed grows exponentially. And the issue here is not just storing the files of significant size but whether the IT system can provide the distributed processing infrastructure to make sense of the genomic data. Not only must the system scale up, but it must be able to scale back down and provide the elastic scalability necessary to be cost effective.

Yet, despite the challenges, obtaining this data is crucially important for advances in healthcare. The ability to provide precise medicine to an individual based on their DNA helps in the identification of risk factors and possible treatments. An example of that is close to home. I experienced first hand how a lack of cognitive support in physician workflow at the point of care nearly saw my son, diagnosed with a rare form of Cystic Fibrosis at the age of 9, nearly missed out on a life-saving clinical drug trial. Fortunately, my involvement in healthcare in the US and passion for technology led me to discover a critical medication for my son when he was 14, which may have been volunteered by the physician we were working with if he had enjoyed access to real-time information about new treatment procedures. Today, my teenage son plays football for his high school team. This is a success story that could be repeated across the nation if the right information could be made available to the right people at the right time.

Thus, vital now are systems that can handle the massive volumes and velocity of data that is and will be created – not just in terms of storing that data but interpreting and analysing the data in real time to Flexibility is also critical, given that all these emerging areas of technology – personal monitoring, risk factor identification, genomic data analysis – are in their infancy and we do not yet know all the risk and social factors associated with them.

What about privacy?

The spotlight is on healthcare data privacy globally, particularly in light of recent database hackings. The FBI warned in August 2014 that healthcare industry companies are being targeted by hackers, following an attack on US hospital group Community Health Systems Inc that resulted in the theft of millions of patient records.[21] Despite this, a higher percentage of people said they are more willing to share their health records (47%) than their phone records (38%) or banking information (30%) to aid innovation.[22] To encourage data sharing and streamline the repetitive nature of granting waivers and data-rights administration, it may be better for data approvals to follow the patient, not the procedure.[23] Privacy requirements will continue to evolve, as big data and access to it expands, but as long as the patient is in control of who can, and cannot, access their record, then personal preferences can be supported, and those that want the best care can enable it.

The other significant trend in privacy and security at the moment is the trend towards Multi Factor Authentication. This trend, in response to local breach tactics, requires not only password protection, but also the

[20] http://www.technologyreview.com/featuredstory/535016/internet-of-dna/

[21] http://www.cnbc.com/id/102398852#.

[22] http://newsroom.intel.com/community/intel_newsroom/blog/2013/12/09/the-world-agrees-technology-inspires-optimism-for-healthcare

[23] 'The 'big data' Revolution in Healthcare'; McKinsey & Company

use of a physical device that the user must have in order to access the system. The most applicable method for the "device" in modern times is the use of a smartphone or SMS technology that requires the user to enter the code immediately following their username/password combination. The multi factor authentication means that, in order to access the system, not only does the hacker require a password, but also requires access to the physical device matching that account.

This significantly reduces the chance of a breach based on stealing identities.

The big data engine

Big Data necessitates technology that can handle it – massive amounts of data, in many different formats, arriving very fast. But the technology needs to do more than consume large volumes of information. Increasing data volumes and complexities in healthcare are pushing the limits of traditional Relational Database capabilities. Complex processing mechanisms are required to extract and aggregate valuable information from raw data feeds and handle structured, semi-structured and unstructured data. The need is for a healthcare focused, big data processing engine that enables real-time risk stratification and the ability to follow up on that risk assessment for a population of millions; a system that can be modified and updated regularly and can use historical data to assess new risk criteria as they become known. We need a solution that addresses scalability, flexibility, data liquidity and fault tolerance.

The information contained in any given health message can serve many different purposes. For example, an ADT message has administrative as well as clinical data. It might update the Patient Demographics, contribute to a Risk Assessment or provide information contributing to point of care decision support. A C-CDA document might contain smoking status thus affecting stratification of a population as well as lab results contributing to the clinical record, perhaps suggesting that the patient would benefit from being assessed for a pre-diabetic profile. Workflow sensitive information can arrive out of order, incomplete messages may require editing and re-submitting, improved analytic algorithms might benefit by replacing data already processed and coded data may be absent, incomplete or using terminologies that the recipient doesn't recognise requiring normalisation.

Therefore, the Big Data Engine must be able to effectively address all of these challenges. It must be able to present data in a range of different data models, for their specific purpose.

For example, medication dispense data might inform Care Gap analysis routines, be part of the longitudinal medical record, contribute towards cost analysis and feed stratification routines that have a medication component. The Big Data Engine must also have the ability to standardise data models, because an explosion of minor variants of data models is confusing rather than enabling. Core, standards-based models that can be extended in a controlled way are required to ensure that data can be easily utilised by consuming systems. Emerging standards such as HL7's FHIR will prove to be extremely valuable for this purpose.[24]

[24] https://geekdoctor.blogspot.co.nz/2013/06/the-june-hit-standards-committee.html

The importance of analytics

Analytics is a key component of Population Health Management and a critical adjunct to the Big Data Engine. It underpins much of the value that data can provide to clinical care and insurer assessments, in the same way that the banking industry globally uses trusted data and analysis to derive insights from billions of transactions for better targeted offers. It is required to measure outcomes and costs for every patient. Analytics is needed to make sense out of raw data but it can be a difficult, iterative process. You start with an algorithm, run it on some data, evaluate the results, alter the algorithm and try again. In doing this, the ability to replay and extract additional content from existing sources is extremely useful.

Big Data analytics can provide stratification and predictive modelling. It can be difficult to identify people in need of medical attention and in particular, the critical timing of intervention to reduce the impact of a chronic illness. Using one of the many models available (for example, the Johns Hopkins ACG System) to group patients into risk layers assists care providers and healthcare organisations to take the actions necessary to achieve improvements in population health.

In addition, a key driver of PHM is to reduce costs, which involves, in part, making better use of resources – both financial and workforce. There are algorithms to measure outcomes and relate those back to the interventions and the cost/effectiveness of those interventions at an individual and a population basis. Big Data can be analysed to provide information for health system managers tracking organisational objectives such as cost reduction and patient satisfaction improvements,[25] on a peer review basis between clinicians looking at different approaches to similar clinical situations; and at the point of care to influence clinician behaviour to achieve equivalent outcomes at the best possible cost.

There are three parts to an effective population health analytics solution:

Store first, analyse as needed. It is impractical to have an understanding of all expected analytical needs up front and establish designs for data models that can support these.

Analyse on demand. As needs for additional analytical information arise, the system must be capable of retrieving previously stored raw data and applying new algorithms to the data models. For example, a clinical trial may benefit from raw data already in existence that was previously cut a different way.

Adapt to local needs. Enable even small groups of users to get the information they need. The tools and methods should support retrieval and conversion of raw data into any-sized pools of information quickly.

Once algorithms are determined, they need to run continuously against the incoming data, updating models in real-time as data is processed.

[25] Merlino, James and Ananth Raman. "Health Care's Service Fanatics: How the Cleveland Clinic Leaped to the Top of the Patient-satisfaction Surveys." Harvard Business Review 91, no. 5 (May 2013):108-116.

Open up to innovation

Transforming healthcare is not just about analysing risk factors and coordinating the care of patients. It is about creating disruptive change that allows innovation to rise to the fore, where the best ideas and "game changers" can be successful. The current closed ecosystem of EMR vendors and systems that retain control of the data within them stifles innovative plays by smaller companies. In 2008, Apple enabled innovation on a massive scale when it launched the App Store, calling for applications targeted at iOS devices that made use of specific attributes of those devices. The App Store model meant developers could build successful software applications without the backing of a large company.[26] Similarly, in healthcare, opening the longitudinal patient record to organisations that would otherwise have no access to the data will open up possibilities. In this environment, many startups would develop niche applications for specific conditions that would not necessarily gather attention from large corporate EMR systems.

The best way to enhance innovation is to create an open platform, where information aggregated from multiple sources can be shared via Open APIs that allow secure, authorised access to discrete information in a patient's record. The JASON report prepared for the US Agency for Healthcare Research and Quality promoted an architecture where healthcare information was more interoperable via a set of Open APIs with the goal to substantially improve the interoperability of healthcare data.[27] Opening up a platform will begin to shift control from a small number of software vendors to a software ecosystem with a diversity of products and "apps" focused on the patient, enabling healthcare providers to partner with patients in data sharing.

In an open system, a patient will have increasing control over his/her own data and will take responsibility for that information by reviewing the elements of the EHR, setting access permissions, and making his/her own contributions to the dataset. Increased patient engagement will foster improved patient education, health maintenance, and treatment compliance. Physicians and other health care providers will become discerning customers of a robust health data infrastructure, rather than slaves to a closed-box system. Patients and providers will gravitate toward user interface applications that provide the best functionality and convenience and serve their particular purposes.

The next question regarding an open system is whether today's interoperability standards are a help or a hindrance to many organisations sharing information. Currently, standards are heavyweight and difficult for an innovative startup to implement. An exciting development in this space is an emerging standard from HL7 called HL7® FHIR® Standard: Fast Healthcare Interoperability Resources (FHIR). FHIR can help with removing the common barriers to interoperability, making information available to those who need it; allowing implementers to focus on the more complex issues of clinical behaviour, rather than the "nuts and bolts" of representing and moving information around; exposing information with widely used open standards, which enables a "clinical ecosystem" where innovators can interact with data that was previously unavailable in proprietary silos of information; and standardising how data is represented, making it straightforward to apply more advanced technologies.[28]

---

[26] http://www.computerworld.com/article/2478691/smartphones/apple-s-biggest-innovation-for-2008---the-iphone-app-store.html

[27] http://healthit.gov/sites/default/files/ptp13-700hhs_white.pdf

[28] Dr David Hay, "FHIR Ignites Healthcare Sharing", Orion Health White Paper, March 2015

The solution

To carry out successful population health management, healthcare organisations need an open enabling information technology platform that has the ability to:

- Store raw data and extract future value from it
- Merge data from all sources – claims, clinical, social and consumer as a start
- Intelligently process that raw data and hold normalised data in stores fit for purpose
- Apply new and innovative analytics to that data
- Share that information; and
- Allow other organisations and individuals to innovate on the data they access.

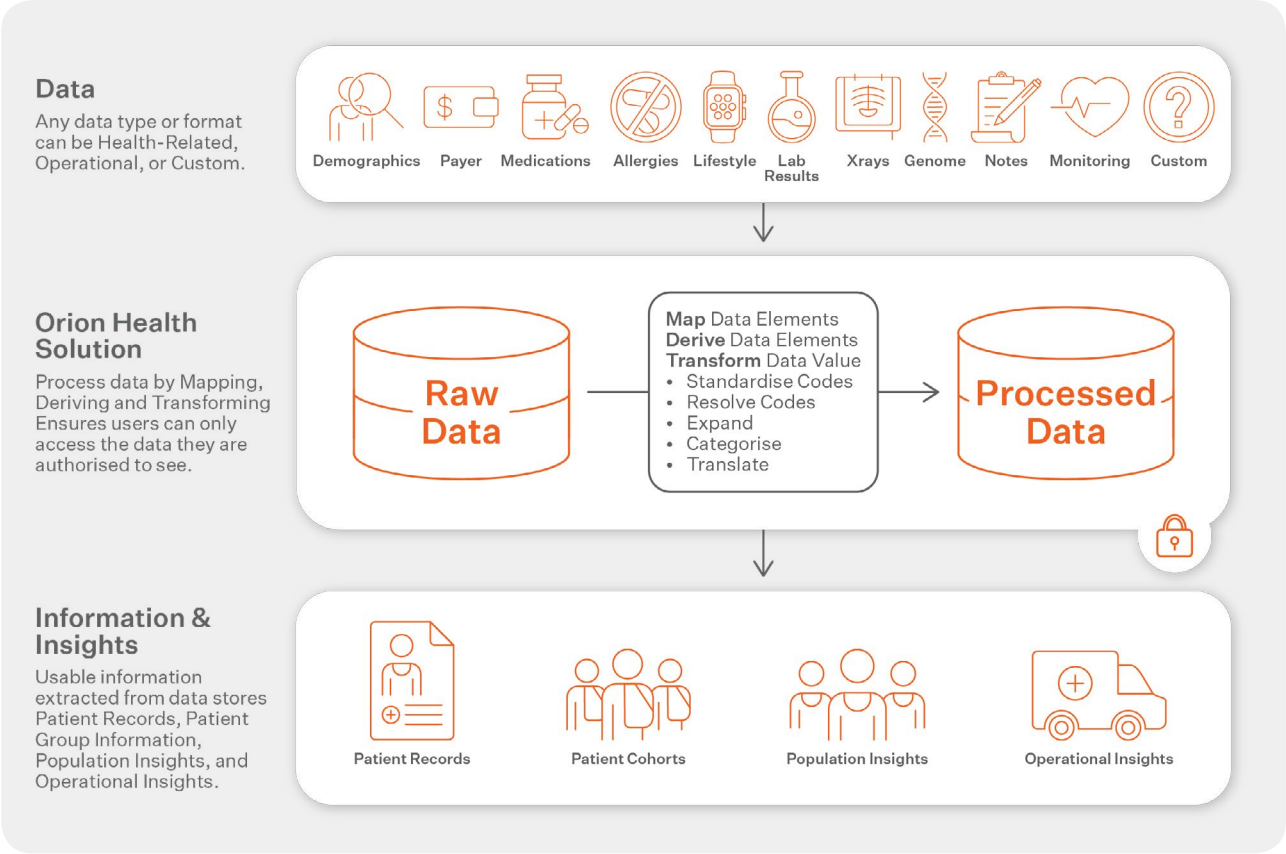

FIGURE 1: Data Processing Pipeline

Orion Health's open platform, powered by Amadeus, is the true big data solution for health, combining all forms of data in one raw data store; processing that data through the Data Processing Pipeline; and storing that normalised data in specific data spaces, ready to be accessed through Open APIs and analysed to gain insights into information.

Our data platform and distributed approach makes us one of the first vendors to enable near real-time services for workflows and measurements. This enables near real-time care management, medication management, and patient engagement. Synchronisation and loop back problems are eliminated, enabling the right care at the right time.

Orion Health open platform is turning Big Data into the big picture, providing an integral tool for managing the health of individuals, communities and populations, globally.