

DirectorTech

Strategy + Risk + Governance

001

Quarterly high-level advice on untangling business technology from the Centre for Internet Safety at the University of Canberra.

Issue 001, September 2017

Welcome to DirectorTech

Nigel Phair, *Managing Editor*

Technology is transforming organisations faster than ever before, so understanding it is vital. But most technology-related news is written by technologists for technologists, and in the rush of the continuous news cycle there's little time to reflect.

DirectorTech intends to address that. We're stepping back to reflect. We'll look back at some of the key technology stories, and dig into them to highlight their lessons for company directors – in the language of company directors.

After all, company directors have a responsibility to inform themselves, so they can make the big decisions about the strategy, risk and governance of technology in their organisation.

Enthusiasm for the opportunities afforded by new technologies has to be balanced by a sober analysis of the potential risks, and that understanding has to adapt to changing circumstance. As our analysis of **Australia's 2016 Census debacle** describes, there were many technological failures, but overarching was a breakdown of project management.

Good technology governance is also vital. That's enabled by the board, working with senior executives to establish strategies, structures, policies and processes for managing the technology. The aim is to improve business process, and therefore business results.

The Victorian Government's Independent Broad-based Anti-corruption Commission's inquiry into the state's **failed Ultranet project** shows what happens when governance breaks down: corruption and \$180 million down the drain.

Moving business processes into the cloud and providing mobile access can mean increased agility and massive cost savings, but it also can also make it harder to detect and resolve security breaches. Meanwhile, the average financial impact of data breaches is increasing. Information security, a task that was once delegated to the IT staff, has now become a strategic issue. Failure can ruin organisations.

We want DirectorTech to keep evolving to bring you what you need.

Please **let us know** what you think, **subscribe now**, or sign up for **free email updates**.



DirectorTech 001

001.001	
Why data governance is so important now	4
001.002	
Virtual Reality: the need for strategy	6
001.003	
Governance and the sharing economy	9
001.004	
The business risks of artificial intelligence	12
001.005	
How Ultranet went ultra wrong	15
001.006	
Lessons from the Censusfail omnishambles	22
001.007	
When algorithms turn evil: the Centrelink debacle	30
001.008	
Plan for your inevitable cybersecurity breach	35
001.A1	
About the Contributors	38
001.A2	
About DirectorTech	39

001.001

Why data governance is so important now

In the first of a series of columns, **Kate Carruthers** introduces the practicalities of treating your data right.



Data governance is an increasingly important discipline for businesses to adopt. The words “data governance” might sound fusty and rule-laden, but it is important, because data is the foundation upon which every digital transformation rests.

Data governance is a practical discipline. It’s about managing data as an asset, and it asks the following important questions:

- Do you know the value of your data?
- Do you know who has access to your data?
- Do you know where your data is?
- Do you know who is protecting your data?
- Do you know how well your data is protected?

Data is a key business asset, so every business needs to know the answers to these questions. This is where data governance comes in. It helps organisations understand who can make decisions and who ought to be consulted in respect of their data.

I use the following definition as an anchor when introducing business executives to data governance:

“Data governance is the organisation and implementation of policies, procedures, structure, roles, and responsibilities which outline and enforce rules of engagement, decision rights, and accountabilities for the effective management of information assets.”

— John Ladley, “Data Governance: How to Design, Deploy and Sustain an Effective Data Governance Program”, 2012.

Increasingly, businesses are moving towards using data and analytics in new ways. Data governance can help to identify where the data is, how it's being managed, and who is accessing it.

Further, the rapid growth of the Internet of Things (IoT) means that businesses are accumulating ever-greater volumes of data, so the challenges of managing that data will also increase. Those challenges will become ever more important for boards and senior executives to understand.

Data governance also feeds into issues like digital ethics and privacy. As we venture further into the world of big data and data analytics, these issues will become more important as data starts to be joined up in new ways.

Next time we will delve a bit deeper into data governance and how it can help.



001.002

Virtual Reality: the need for strategy

*Businesses are already adopting VR, but **Nigel Phair** wonders where the strategy is.*



Virtual reality (VR) is, according to the **Wikipedia entry**, technology which can “generate realistic images, sounds and other sensations that simulates a user’s physical presence in a virtual or imaginary environment”.

VR isn’t just about online gaming. It’s already creeping into many aspects of business and social life. If VR becomes as pervasive as its enthusiasts believe, any organisation without a virtual reality strategy will soon be left behind.



Virtual reality’s arrival in marketing

Big companies have already taken advantage of virtual reality, especially in their marketing space.

In 2015, Marriott Hotels created an experience which **“teleports” guests to different places** so they can experience the destinations through their senses. The experience included visual stimulation through a headset, as well as the physical sensations of heat, wind and mist. **Qantas has a similar setup** in some of their airport lounges.

Also in 2015, Hyundai launched a **4D Hyundai driving experience** based on Oculus Rift technology. This multi-sensory experience combined a VR headset with a moving platform which added motion, and a soundtrack complete with road noises.

These experiments were designed to draw in consumers, and encourage them to become part of an experience, rather than be “sold to” through regular marketing and advertising.



Virtual reality in small organisations

Virtual reality doesn't have to be as big and complex as those two examples. Consider the recent **Pokémon Go** craze, where people were catching Pokémon just using their smartphones. Small businesses were able to trade on this virtual experience by registering as PokeStops to lure potential customers inside.

For small organisations, VR means being able to conduct virtual meetings, or offer product walk-throughs to potential clients. It means being able to test a product during development, finding faults before investing in production. It also offers a way to train people on specialised equipment or techniques without exposing them to injury, or risk damaging expensive equipment.

As with any other business decision, however, you need to show that implementing VR into any aspect of your business will provide a clear value to your organisation.



Developing a virtual reality strategy

Your strategy must consider how VR can increase your organisation's competitive advantage while still meeting the requirements of your vision and values.

- How are you going ensure your organisation stays up to date with advances in virtual reality?
- What might work and what might not?
- Where will we get the most benefit from VR?
- What is the most cost-effective use of VR?
- What is the likely return on investment?
- Who is responsible for managing the VR experience?
- Is there another way to achieve the same result without yet having to tap into this technology?
- How will you know when your organisation is ready to adopt VR?
- What resources do you currently have available?

Part of your strategy development should also include an analysis of your own role as company director. How and where will you fit into the overall strategy, and how will you stay up to date with current developments in VR?

At the moment, there appears to be very little information available on developing a virtual reality strategy, but we do know that consumers are keen to explore the virtual reality experiences you can offer.


Your challenge will be to develop a strategy that keeps you up to date with technological advancements and one step ahead of your consumers.



001.003

Governance and the sharing economy

As **Nigel Phair** notes, the disruption of existing business model can cause confusion.



The sharing economy is thriving, thanks to the arrival of smartphone apps, social media, and a tightening economy. Yet the growth of the sharing economy has brought with it confusion over rules, rights and responsibilities.

“The sharing economy [is] an economic model in which individuals are able to borrow or rent assets owned by someone else.”

– Henrique Schneider, “Creative Destruction and the Sharing Economy: Uber as Disruptive Innovation”, 2017.

The sharing economy model is most likely to be used when the price of a particular asset is high and the asset is not fully utilised all the time.

Airbnb and Uber are two of the best-known examples of the sharing economy. Other examples growing in number are co-working spaces, shared equipment, shared technology, shared transport, and even the possibility of shared staff.

In 2016, **PricewaterhouseCoopers predicted** that transactions in Europe’s five most prominent sharing economy sectors – collaborative finance, peer-to-peer accommodation, peer-to-peer transportation, on-demand household services and on-demand professional services – could see a 20-fold increase to €570 billion by 2025.

That’s just Europe.

It makes good business sense to consider tapping into that potential income.



The governance implications of a sharing economy

While the sharing model often makes economic sense, it raises questions about how to manage risks, security, and even safety. When you don't own or even co-own the shared item, where does your responsibility begin and end?

The nature of governance is changing. The sharing economy is changing traditional two-party relationships into relationships with multiple parties, all with different rights, obligations, and expectations.

Uber had to deal with all these questions during its set-up, and continue to re-visit them as the company expanded — sometimes controversially.

- Is Uber a taxi service? In February 2017, the Federal Court of Australia ruled that **ride-sharing is a taxi service**. That classification determined the amount of tax its drivers must pay.
- How do you manage individual performance under a sharing economy?
- Is Uber legal in all states?
- Are Uber's drivers regarded as employees or independent contractors, and how does that affect its employment responsibilities?
- Are Uber drivers entitled to Workers Compensation?
- Who is responsible for public liability if something goes wrong and there's an accident or injury?

Uber's directors have had to battle over and redefine almost every aspect of their business, despite the simple concept it was intended to offer.



Exposing governance gaps

Current regulatory frameworks were established when there was a clear ownership of the property or service. Under a sharing economy, there are no such clear boundaries.

If you are a director of an organisation entering the sharing economy space, you should understand that there may be no precedents for you to follow.

- How will you manage the sharing of tangible and intangible assets?
- How do you measure and value assets such as time and money?

- How will you manage and monitor the use of tools, stationery, equipment ,and consumable items?
- How do you safeguard your employees, if you have them?
- How do you manage equity between the sharing businesses?
- How do you keep abreast of regulatory changes affecting your business?
- How will you maintain the “peer-to-peer” approach that has been so important in establishing the sharing economy?
- How will you use and manage social media to support your sharing economy?
- How will you insure yourself, your business, your service and your team?

Under a traditional business model, governance requirements are clear. The new sharing economy model means sending directors back to the beginning to analyse every aspect of their organisation.


While the future may appear to be filled with nothing but questions, it’s actually a time of new opportunity. The first organisations to tap into the sharing economy may be able to secure a wonderful competitive advantage, and seize the opportunity for growth.



001.004

The business risks of artificial intelligence

*If you trust a machine to make your decisions, you'd better be prepared, writes **Nigel Phair**.*



Artificial Intelligence (AI) is the science and engineering of making intelligent machines, especially intelligent computer programs. Many movies have sensationalised the topic, but AI is real, and it's already part of our lives.

How prepared are we for the possible consequences of relying on AI to make our decisions for us? Because to a certain extent, that's just what AI does.

As a director, you need to ask yourself some tough questions about the risks this poses to your company.



What is AI?

Artificial intelligence is a broad term, covering many approaches to making computers mimic an ability to work on problems normally requiring human intellect. How humans think. A key goal of AI is to develop systems capable of operating, at least to some degree, independently of human control.

For example, DeepMind, owned by Google's parent company, Alphabet, has taken AI a step further by developing a neural network which can not only learn, but remember and recall facts. The system can then make choices based on memory and logic.

In a paper published in October 2016, DeepMind reported that the system could **navigate through London's Underground** without human input.



Making mistakes faster than ever before

The biggest risks associated with AI are:

- Although it may speed up our processes and increase productivity, it will also therefore speed up the negatives associated with business, and expose or magnify errors and losses.
- There's the potential for mismanagement of AI by human agents
- There exists the possibility of AI malfunction with negative consequences

An added complication is that the systems are new, and continually evolving, so it's almost impossible for computer engineers to keep up with understanding and managing the risks.



What might an AI that can make its own decisions do to your business?

Let's not lose sight of the humans. Many of your clients will be worried by the concept of AI, knowing little about it other than the Hollywood concepts. There's every possibility they won't accept the non-human contact. While AI like Apple's Siri captured public imagination, we also know that robot voices at the end of the telephone support line are unpopular.



How AI will challenge the board

Boards may be comfortable with traditional areas such as human resources or finance, but few directors are as comfortable dealing with technology. Understanding AI systems is yet another challenge, as is being prepared for possible attacks or malfunction without even knowing how this might happen.

How can a board manage a risk when they don't understand enough about it to pinpoint danger zones?

The other question is about who holds the ultimate responsibility for AI.

"Paralysis occurs because AI does not fit neatly into any specific area of responsibility; it is not exclusively for the chief information officer, the chief marketing officer or even the chief operating officer."

— **Josh Sutton**, global head of AI practice. Publicis.Sapient

Organisations which are “agile” can adapt quickly, yet a 2015 McKinsey and Co survey of 1000 organisations showed that **only 20 percent were truly agile**.

Perhaps the board’s focus should be on producing agile companies, which can incorporate and take advantage of AI and technology, rather than focusing on the technology itself.



Preparing directors for AI

In 2014, Japanese venture capital firm Deep Knowledge Ventures responded to the AI challenge **appointing an algorithm its board of directors**. The AI, named Vital, was tasked to sift through large amounts of data to identify the best vehicles for VC funding.

While that could be the way of the future, plenty of groundwork needs to be done first.

To minimise risk, directors need to consider the composition and structure of the board itself, to ensure it has the skills and competency to deal with AI and other technological advances.

The application of AI is broad, so it may take two or three specialist directors to fill the knowledge gap.

It’s important to understand , however, that these directors must be involved in the regular work of the board, so their expertise is applied at the right time and place. By working together, the directors can minimise the risk to people, jobs and systems.

Never forget that AI must integrate with the overall business strategy of the organisation, rather than interfere with it. The responsibility will fall to the directors to develop a sound strategy to incorporate AI, so it enhances your ability to reach the overall goals.

The strategy must also include a risk management plan, to identify and minimise risk.

It may be that the board needs to create a sub-committee with a special focus on technology and AI. This will enable the board to be alerted to critical issues early enough to act upon, and to help embed AI into the organisation’s processes with minimal disruption.

However you approach the topic, it’s critical that every board member is able to understand the possible impact of AI on your organisation, its productivity and its market competitiveness.


Without this, the board will be ineffective, unable to pinpoint the questions, let alone answer them.



001.005

How Ultranet went ultra wrong

*Victoria's ambitious education network was a spectacular failure. As **Justin Warren** explains, it badly managed and riddled with corruption.*



In January 2017, the Independent Broad-based Anti-corruption Commission (IBAC) published its **Operation Dunham special report** into Victoria's Ultranet project. The report is well-written, clear, and thorough in its treatment of the issues uncovered by IBAC's investigation, and it is mercifully free of the usual weasel-words and buzzwords that clutter too much of today's writing.

Ultranet was to have been an online teaching and learning system for all Victorian government schools. It was announced in 2006 by then Premier Steve Bracks, and the programme — to be run by the Department of Education and Training — was given a budget of \$60.5 million.

The Victorian Auditor-General's Office (VAGO) audited the project. Its **report**, published in December 2012, found that "Use of the Ultranet is well below expectations, with only 10 percent of students and 27 percent of teachers logging into the system."

In 2013, the project was closed down after having blown anywhere from \$127 million to as much as \$240 million.

The report has clear lessons for the governance of technology projects generally, though few of them are new.



Systemic failure

IBAC's investigation found that the tender processes for the Ultranet were corrupted due to improper relationships between senior officers of the Department, primarily with Oracle Corporation Australia Pty Ltd (Oracle), and then with CSG Services Pty Ltd (CSG).

While there were multiple, serious incidents of corruption uncovered by IBAC during its investigation — including senior officers potentially breaching insider trading laws — I choose to focus here on the systemic nature of the issues uncovered, rather than individual instances.

As IBAC stated in its report, “it was the collective failure of the Department’s three ‘lines of defence’ that ultimately allowed the conduct under investigation to continue unabated.”

The first line of defence was the reliance upon “the ability of managers and leaders to follow correct procedures and to act with integrity.” Multiple senior officers in the department failed to act with integrity. Darrell Fraser, former Deputy Secretary, in particular, went to extraordinary lengths to ensure the CSG/Oracle offering was successful, frequently ignoring the correct procedure.

Yet it is the failure of the other lines of defence that make clear that the Department suffered a systemic failure. Individual components of a system can, and will, fail. Organisational resilience requires other systems to provide checks and balances, reinforcing each other to protect the organisation as a whole.

The second line of defence in the Department was “the various systems and processes that are in place to safeguard and regulate activities performed by individuals in the first line. This includes financial approval processes, procurement processes, and governance committees and frameworks.”

The regulation of the actions of individuals failed. In some cases multiple individuals were colluding to circumvent controls, with one approving the actions of another. In other cases, processes were circumvented by bad actors providing false information that was not independently verified by those responsible for approvals.

These issues should have been caught by a third line of defence, “the assurance provided by the Department’s audit functions that the systems and processes specified in the second line are working appropriately.”

The Ultranet project avoided audit scrutiny because whenever it was to be placed on the Department’s audit plan, Rosewarne and Fraser would push back. Gateway reviews, “conducted by an independent review team at critical points of a project’s lifecycle”, flagged “numerous urgent and critical issues”, yet the project was permitted to proceed anyway.

If any one of these three “lines of defence” had worked correctly, the Ultranet project may well not have been the unmitigated disaster it turned out to be. That they all failed together is typical of a systemic failure.

The causes of this systemic failure were not isolated to individual bad actors, but a result of the culture within the Department. During **Operation Ord**, IBAC exposed “serious and systemic corruption within the Department of Education and Training”.

“The failure to address Mr Fraser’s behaviour time and time again can only be described as a serious failure on the part of some of the most senior leadership within the Department. Operation Ord also exposed this concerning culture among a group of influential senior executive officers at the Department. This same culture made Mr Fraser’s behaviour permissible.”

The failures in the Ultranet project were a consequence of these same, cultural problems.

Systemic failures are indicative of widespread issues that require an holistic approach to fixing them; targeted interventions will not be sufficient. Changes will be required to the organisation as a whole, and therefore require the involvement of the Board and senior executives. It is far too easy to point the finger of blame at individual bad actors. The nature of systemic failures is that it does not matter who the bad actors are, merely that there are some, and there will always be some, hopefully few in number.

Any system that requires all of its components to operate flawlessly, in all conditions, forever, is fatally brittle and doomed to failure sooner or later.



Regulatory capture

The Ultranet case illustrates the risk of becoming too cosy with any one vendor.

Oracle was given preferential treatment in its application to Ultranet. There were obvious conflicts of interest for influential members of the Ultranet Board. Fraser’s inappropriate relationship with Oracle was well-known within the Department.

There is an argument for going with a known and trusted partner, rather than an unknown, but single-sourcing is a two-edged sword.

A strategic alliance can reduce complexity, which reduces costs and time-to-market, but it adds substantial switching costs. Relying on a single vendor reduces negotiating leverage on future deals. It also adds a hurdle if a better alternative arises from a different vendor, reducing your ability to adapt to change.

It can also lead to a de-skilling of the organisation, in favour of relying on the vendor.

As we saw with the Australian Bureau of Statistics’ **Census 2016 failure**, this loss of skill can mean the organisation is no longer able to accurately assess the abilities of the vendor.

Unable to tell if the vendor is performing good work, or providing value for money, an organisation can become dependent on the vendor, and cedes control over its own destiny.

The cost of changing your mind needs to be factored into decision making. The purpose of “agile” methods that are much talked about is to reduce the cost of change; mistakes become less costly and are easier to correct.

A major problem with the Ultranet was the substantial financial and emotional investment by the principal actors, and the **sunk-cost fallacy**, where humans tend to double-down on poor choices in order to “save” the money that has already been spent.

Multi-sourcing, particularly using open standards, provides flexibility to the organisation if it needs to change course, and provides a ready Plan B should the initial direction prove incorrect.

When there uncertainty about the correct which is true of most IT systems as history should make clear — then the degree of flexibility needs to match the degree of uncertainty.



What could possibly go wrong?

The novelty and excitement of new things is particularly prevalent in the technology industry. This unbridled enthusiasm to embrace the new can cause people to overlook problems that are obvious in retrospect.

Why wait until a major project has failed to determine why it failed? A useful technique is to conduct a pre-mortem.

The key to a pre-mortem is to work out all the things that could go wrong ahead of time, and determine what, if anything, can be done about them. By spending time working out the potential flaws in the project, the team can alter the design to avoid the problems identified, clarify their objectives, or ensure appropriate safeguards are put in place to manage issues as they arise.

Pre-mortems differ from critical reviews in important ways: critical review is performed after the fact, by an external party. A pre-mortem is performed as part of the design process by the team itself.

Criticism by an external party tends to create an adversarial relationship, where criticism of the project is taken personally by the team who has invested so much of themselves in the work. A pre-mortem can help to alter this antagonistic relationship to one of colleagues

working on an interesting problem and solving it. Making a system resilient to failure becomes an explicit design objective, rather than a process of avoiding scrutiny.

Assessing potential risks and guarding against them ahead of time is smart preparation. If a team is planning to scale Mount Everest, and someone points out that ordinary shoes are not a good choice, that is helpful advice, not pointless negativity. CXultural norms should be set by senior leadership that detecting flaws — and helping to fix them before they become disasters — is every bit as important as creating new things.

The pre-mortem should also include the risks identified to the project. “What if the risk assessment turns out to be wrong?” is an important question to ask.



Who watches the watchers?

“Quis custodiet ipsos custodes?”

— Juvenal, “Satires”

There is a remarkable ability of those designing systems to believe in the inherent goodness of humankind, given the abundance of evidence to the contrary.

It is important to strike a balance between trusting people, and checking to see if they make mistakes. Most errors are innocent, but some are not. In general, the best system design is one that makes it easy to do the right thing, and hard to do the wrong thing.

With Ultranet, bad actors were able to bypass checks and balances due to systemic cultural issues within the Department. Leadership from the most senior levels of the organisation is required to show not only that the rules must be followed, they must also be seen to be followed.

When designing a system, too often people believe it best to err on the side of caution, and concentrate on making it difficult to do the wrong thing, with the side-effect of making it difficult to do the right thing as well.

Such a system has the opposite of the intended effect. People will bypass the system so they can get their work done, undermining the whole approach. The response is generally to erect ever-higher barriers, like stern threats of dismissal if people violate company policy.

Have you read all of your company's policy documents? Have you followed all of them to the letter on all occasions?

It's easy to justify disobeying a system that people consider ridiculous, or unnecessarily burdensome. The hierarchical nature of organisations tends to allow more senior people

more leeway over which systems they disobey. We saw this with Ultranet, and yet senior people have greater scope than junior line-staff to damage the organisation.

With great power comes great responsibility.

Finally, an important question to ask is “What if someone simply doesn’t follow these rules?” Relying on individual integrity, as we see with Ultranet, is not enough. It is important to ask not “Who will let them?” but “Who will stop them?”



Slow down to speed up

The technology industry is currently infected with the idea that everything must move at breakneck speed, and anything that gets in the way is of little value. Security and compliance are seen as impediments.

In some cases this is true, in the same way that your immune system is an impediment to a lethal infection. A disruption to your heart's ability to pump blood and oxygen to your brain is not a good thing.

The potential advantages from IT systems are great, this is true, but only if those systems make it into production on-time and on-budget. Projects that fail waste time and money that could have been spent on other, useful systems that will now never get the chance.

How many successful initiatives will we never see in the constantly resource-constrained education sector because Ultranet deprived them of \$240 million of sorely needed capital? That’s a lot of books and pencils.

IT projects are like any other investment, and should be managed as such.



Your Next Steps

- Encourage teams to perform a "pre-mortem" as part of setting up any project. What resources will be needed to overcome issues if they happen? What warning signs can you look for to detect if one of the risks is about to actually happen?
- Do a pre-mortem of your own governance systems. If a bad actor were to attempt to defraud the company, how would they do it? Would you be able to notice before major damage is done?

- Evaluate your vendor relationships. Are you too close to particular vendors? Why have you aligned your business with certain partners, and are those reasons still applicable? What will it cost you to change your mind?
- Think critically about your own behaviour. Do you provide a good role model for the correct behaviour in your organisation? Ask those who work closely with you the same question and see if they agree. You may be surprised by what you hear.



001.006

Lessons from the Censusfail omnishambles

*Poor management helps cause most of the Australian government's recent string of technology failures. **Stilgherrian** says that's often down to cultural issues.*



When it comes to managing big IT projects, the Australian government has been showering itself in something other than glory. Since the middle of 2016 we've seen three highly-visible train wrecks, and a plethora of slightly more minor disasters.

In November 2016, the Australian Taxation Office (ATO) suffered a systems outage lasting a couple of hours. It was merely the precursor to a series of major outages, some lasting for days, that continued through to February 2017. It's still not fixed.

From December 2016, it became obvious that Centrelink's automated debt identification and recovery system was a mess — organisationally, socially, and ethically. It has inflicted distress and financial pressures on the very people that Centrelink is meant to be serving, and it's unlikely to achieve its budget goals. But the warning signs were there at least a year earlier. If only someone had known to look.

But before those two came the most public failure of all, and that's today's case study.

In August 2016, the Australian Bureau of Statistics (ABS) had planned to run the 2016 Census online. But their systems collapsed under the load of a series of distributed denial of service (DDoS) attacks. The eCensus became Censusfail. Humiliated, the ABS was forced to shut it all down.

All three cases have something in common. They didn't involve a failure of technology, but a failure to manage the design and implementation of technology, and a failure to understand and to mitigate the risks.

Boards and executive management teams need to get this stuff right.



The eCensus dream that became a nightmare

According to the ABS, the Census is “Australia’s largest peacetime logistical operation”. It’s a long-trusted exercise involving 40,000 employees, 20 million postal items, and two million phone calls. Choosing to move all this data collection online was a no-brainer.

The ABS had already run an online system for the 2011 Census, but in 2016 the eCensus would become the primary mechanism. Paper was to be a last resort.

And so a massive communication campaign urged Australians to go digital, and fill in the eCensus on Census night, Tuesday 9 August. And so they did. Until it fell over.

If schadenfreude is to your taste, then the two government inquiries into Censusfail will feed even the biggest gluttons.

- **2016 Census: issues of trust** by the Senate Economics References Committee.
- More readably, and my personal recommendation, the **Review of the Events Surrounding the 2016 eCensus** (PDF) by Alastair MacGibbon, the Prime Minister’s Special Advisor on Cyber Security.

MacGibbon’s 93-page report includes a 15-page timeline. It details a cavalcade of technical failures, communication failures, and poor crisis management.

Read it in full, if you have the time. Meanwhile, here are the highlights.

- Mid-morning, hours before Census night’s post-dinner peak times, the eCensus system started failing under DDoS attacks. This wasn’t a flood of malicious traffic, though. It was a trickle, less than one percent of the volume that other organisations deflect every day.
- The ABS and their prime contractor IBM implemented their agreed mitigation procedure. It worked at first, but a third of the malicious traffic was still getting through to the eCensus system, which couldn’t cope with even this small extra load.
- The Australian Signals Directorate (ASD), the nation’s cyber defence organisation, was called in for technical advice. That advice was incorporated into the increasingly ad hoc and rapidly failing crisis plan.
- When the evening peak period began, the eCensus system simply couldn’t keep up, even though legitimate traffic was within the ABS’s expectations. At one point, a router had to be rebooted, but somehow it lost its settings, and had to be reconfigured manually. The eCensus was overwhelmed.

- Social media was flooded with citizens complaining they couldn't complete the eCensus. But the ABS social media team continued to tell them that everything was working "smoothly as expected".
- While all this was happening, IBM's network monitoring detected an unexplained spike in *outbound* traffic. Could it be malicious activity, perhaps someone stealing the Census data? With that fear on top of all the failures, at 9.15pm the ABS decided to shut it all down.
- The eCensus wasn't back online until 1 day 18 hours and 44 minutes later. All up, the outages totalled 43 hours.



eCensus failures on the night, there were many

- There was no clearly identified and tested **cybersecurity incident response** procedure. That led to messy ad hoc decision making, and improvised fixes.
- Crisis communications were inadequate. It took too long to inform Ministers, stakeholders, and the public. There was no contingency plan or draft talking points for adverse events.

Citizens trying to do their eCensus saw an error message, and it was wrong:

"Please be advised that the 2016 Census online form is currently experiencing high volumes. Please try again in 15 minutes."

- ASD didn't provide timely notifications to the Australian Cyber Security Centre (ACSC), nor to the Department of Prime Minister and Cabinet.
- The government's Cyber Incident Management Arrangements (CIMA) simply weren't up to the task.

As you might expect, all these failures on Census night stemmed from failures in the planning stages.



eCensus cybersecurity planning failures, we've got a few

- There were gaps in cybersecurity arrangements. There wasn't a comprehensive security framework, and there wasn't an independent security assessment.

The ABS's Request for Tender had specified an independent assessment, to be done by an Information Security Registered Assessor Program (IRAP) assessor. IBM agreed to do this, but it never happened.

- The ABS didn't have a formal process for accepting responsibility for system security, including identifying and accepting any residual risks.
- The DDoS protection was completely inadequate, and hadn't been tested. My understanding is that an independent test and assessment would have revealed the flaws well in advance.
- Exchanges between ABS, IBM, and ASD suggest a lack of clarity in capacity, roles, and responsibilities.



eCensus procurement, contracting, and governance failures, we've got a few

MacGibbon lists a dozen issues here, but nearly all of them relate to the “cosy relationship” that existed between ABS and IBM.

ABS is highly dependent on IBM's technology. Moving to another vendor was seen as expensive, and difficult to achieve in what they saw as tight timelines. As a result, the ABS hadn't conducted any market testing or an open approach to the market since 2008

This “vendor lock-in” can cause problems.

- With no open approach to market, how can an organisation be sure they're getting the best solutions or the best value for money?
- ABS held the “unrealistic assumption” that a supplier who had performed well in the past would continue to perform well in the future, even when the scope of services had changed.
- Because the ABS trusted IBM, their scrutiny and independent assessment of IBM's security solutions were inadequate.
- ABS did not have an effective IBM outsourcing oversight framework. It appears that ABS had such long-standing trust in IBM that their assurances were taken at face value.

There seems to have been a complete breakdown of the risk management process. As just one example, MacGibbon explained how the process worked for the DDoS protection.

The 2016 online Census Risk Management Plan – dual-badged to the ABS and IBM, but owned by IBM – specifically identifies DDoS risks and associated planning. Loss of system availability via a ‘technical’ or distributed denial of service attack is identified as ‘possible’ in likelihood, ‘major’ in consequence and ‘high’ exposure

Risk mitigations were put in place, which reduced the risk assessment to ‘unlikely’ in likelihood, ‘major’ in consequence and ‘medium’ in exposure.

In both these instances of risk planning, the high initial risk ratings do not appear to have driven a resultant focus on the effectiveness of implementation of the identified controls. Nor do they appear to have shaped preparedness for incident management and potential supporting communications strategies on Census night.

In short, if a risk was still possible and it could cause major consequences, why wasn’t that planned for?

In October 2016, I wrote that Censusfail was an **omnishambles of fabulous proportions**.

“ABS failed to make sure its contractors were doing their job. As we’ve said so many times, you can outsource work, but you can’t outsource responsibility. Yet the ABS submission to the inquiry is studded with claims that they’d received assurances which they didn’t cross-check.”



eCensus privacy planning failures, we’ve got a few

Starting with the 2016 Census, the ABS planned to extend the time it holds on to the name and address information linked to each census record, from 18 months to four years. But the ABS mismanaged the entire process.

The Privacy Impact Assessment process was limited, with little community input outside stakeholder organisations – a problem, given the added privacy concerns of an eCensus. The ABS failed to explain their justification for this extension to the public. A grassroots protest began, which soon evolved into an ongoing media story, with politicians adding leverage.

The ABS response, however, was to tell those with concerns that they were in the minority and not to worry. As MacGibbon put it:

This coverage created overwhelming ‘noise’ making it difficult for the ABS to remain on message.

The ABS’s planned communications were being drowned out. But rather than trying to adapt its approach to limit the impact the reporting had on the public sentiment toward the Census, the ABS stuck to planned messaging ignoring the public relations storm brewing around them.

The failings of the ABS to address issues of concern in the media extend to its use of social media. Analysis conducted on ABS Twitter and Facebook accounts shows that at no point did the ABS significantly change its planned posting schedule or content as a result of critical media reporting .

The ABS did have a social media strategy, of sorts, but it wasn’t flexible. It simply couldn’t react to the groundswell of concern.

- The ABS’s ‘qualifiers’ (thresholds that had to be met to raise concern) were too high. A ‘red level scenario,’ the highest categorisation for negative conversation, was enacted only if someone had 10,000 plus followers or a post had over 30 engagements.
- The ABS’s response to a ‘red scenario’ was simply to hold all social media communications.

“While the ABS did eventually start engaging in the mainstream media, it was too little, too late,” MacGibbon wrote.

If you ever need a case study in how not to communicate a potentially controversial policy, this is it.

The ABS made one fundamental mistake. They continued to talk about the Census and privacy the same way they’d always done. The ABS and the Census were trusted, and the government will take good care of your personal data.

But citizens are now seeing news stories about nation-state digital surveillance, and data breaches, including some by governments. Citizens have learned first-hand how social networks and online advertising companies collect and use their data, and increasingly they don’t like it.

Citizens have also learned how to use social media to organise themselves, amplify their message, and effect change.

The ABS simply failed to adapt to these changing times.



ABS organisational culture failures, we've got a few

"The ABS's business model is old, outdated and in need of renewal. The ABS is almost missing the potential of the digital age by clinging to past practice." MacGibbon wrote, citing an ABS-initiated review of stakeholder relationship health published in June 2016.

Consultancy firm CapDA reviewed the ABS's ICT capacity and capability to conduct the 2016 Census, with a final report in May 2014 . It wasn't exactly positive.

- Although the ABS has project management expertise, rigorous project management is not strongly embedded within the culture and behaviours of the ABS.
- The way the ABS uses agile software development method means that dealing with security, high performance and accessibility are considered late in the cycle.
- It was unclear where, and in whom, the responsibility and authority is vested for making key architectural decisions.
- There was no evidence that any application or data centre performance monitoring is in place.

To judge by subsequent reports, things haven't improved much since.



If in doubt, blame the vendor

In the wake of Censusfail came a barrage of finger-pointing. The ABS blamed IBM. IBM blamed contractors. Contractors pointed to advice not taken, or options not purchased.

In their submission to the Senate inquiry, the **ABS blamed IBM for the failure**. The head of the ABC, Australian Statistician David Kalisch, was the captain who wasn't taking responsibility for navigating his ship of stats.

Two months after that, when the reports came out, Kalisch had changed his tune.

"The ABS underestimated the nature, complexity and risk of the change process," **Kalisch told The Mandarin**. Part of it was organisational pride: the ABS didn't ask for enough help, because they wanted to do it all themselves.

Kalisch also acknowledged that the ABS had failed to adapt. "We expected the public, media and politicians to respond like they had in the previous censuses, and they didn't," he said.



The risk of outsourced IT, that's the big one

As I wrote at ZDNet in November 2016, the two Census reports **highlight government IT incompetence**. Section 6.70 of the Senate report quotes the CapDA review:

"CapDA's report highlights the professionalism and dedication of the staff at the ABS, but in the end recommends that the ABS did not have the internal capacity to develop and deploy an eCensus. If they did not have the ability to develop a solution themselves, it stands to reason that they would only have a limited capacity to question and challenge a contractor employed to develop such a solution."

Exactly. And this is the real lesson.

If no one in your organisation understands the technology, then you can't manage technology projects, whether they're executed by internal staff or external contractors.




Additional Resources

- Full video: **David Kalisch and Alastair MacGibbon on Census 2016**.

When algorithms turn evil: the Centrelink debacle

*Computer algorithm and big data aren't magic, says **Stilgherrian**. They simply implement their creator's view of the world, whether it's right or wrong. That makes using them a matter of ethics.*



Computers are already replacing middle managers. When that happens, organisations become a kind of machine-human hybrid. That's a worry, according to publisher Tim O'Reilly, founder of O'Reilly Media, the publisher of books from which so many network engineers and internet software developers learned their craft.

"When you think about a system like [ride-hailing services] Uber or Lyft, there are programmers who are the top managers, and then there are their programs, the algorithms who are the middle managers, and then there are humans who are actually workers reporting to the algorithms," O'Reilly told the Asia Pacific Regional Internet Conference on Operational Technologies (APRICOT) in Ho Chi Minh City, Vietnam, in February.

"When you have that Uber app, it tells you where to go, who to pick up, passengers rate you, your work scores are reported, and you can be dismissed on the basis of the data that's tracked about you ... Uber has increasingly treated their workers as a disposable commodity, rather than figuring out how to empower them," he said.

"We have algorithms that are increasingly managing our businesses and managing our society. We have to make the right choices to design those algorithms in a way that leads to a better workplace for human beings."

And, it must be said, a better world for their customers, not just for the shareholders.

Uber's prices reflect demand, and during massive demand spikes their so-called surge pricing can be several times the base price. Surge pricing is seen during bad weather, special event or celebrations like New Year's Eve, when public transport systems fail, or even during disasters.

During the 2014 hostage siege at Sydney's Lindt Cafe, Uber's algorithm responded dutifully to the demand of people fleeing the CBD, **raising prices fourfold** and setting the minimum

price for any journey at \$100 . One potential customer's quote for the short run from the Sydney CBD to airport was "\$145 to \$184".

It was more than three hours before a human finally overruled the computer. Customers were refunded, but the brand damage had already been done.

There have even been allegations that some of Uber's drivers have been gaming the system to produce **fake demand spikes**, bumping up the price and therefore their own income. Uber says, though, that they'd soon be detected and sacked. By the algorithms.

Uber is certainly the poster child for algorithmic business, for both good and bad, but it's not the only one. There's plenty of competition to be the most innovative, or perhaps the most extreme.

Some supermarkets are introducing digital pricing displays so that prices can be changed quickly and easily, as the BBC reported under a delightful headline, **Why your bananas could soon cost more in the afternoon.**

"Marks and Spencer conducted an electronic pricing trial where sandwiches were sold at a discount in the morning to encourage shoppers to buy their lunch early.

While the company isn't currently planning to do this more widely, it comes as several of the large supermarkets are trialling the idea.

Sainsbury's says it ran an electronic pricing trial two years ago — but wouldn't say what conclusions it reached — while Morrisons and Tesco are each currently trying out the system in one of their stores."

This kind of demand pricing might provide bargains and increase sales, but it might also lead to price-gouging, and disadvantage customers who might have no choice but to shop at high-demand times.



Centrelink, just following (the computer's) orders

In mid-2015, Centrelink, the Australian government's social security agency, started using algorithms to help recover overpayments, such as when someone continued to be paid Newstart benefits or Youth Allowance after they were no longer eligible.

The plan was to cross-match Centrelink's data with data from the Australian Taxation Office (ATO) to identify discrepancies.

Centrelink sees these overpayments as a debt to be recovered. The system was therefore designed to calculate the amount owing, and begin the debt recovery process. The government claimed the system would recover \$4 billion.

The program has been, and continues to be, a shambles.

The algorithms were so poorly conceived that vast numbers of incorrect debt notices were issued. The subsequent procedures for Centrelink clients to dispute the debt were onerous, with tight deadlines, and the amounts claimed were often impossibly high, causing great distress.

Details of the problems and their scale are still being winkled out of the Department of Human Services through freedom of information requests, but here's just a small sample of the problems.

- Centrelink tracks income fortnightly, but the ATO works to annual cycles. Income can vary through the year, so ATO data isn't sufficient to calculate eligibility.
- Once debt notices were sent, any challenge had to be lodged within 21 days, and it could only be done online.
- Centrelink was chasing overpayments from up to seven years ago. So a challenge might require employment records going back that far, even though the ATO only requires records to be kept for five years. Centrelink itself had told clients they only needed to keep payslips for six months.
- Centrelink's process for cross-checking the employer details provided in a challenge was severely flawed. Consultant Justin Warren, who has analysed that process, described it as **mind-bendingly stupid**.

The system failed to match the word "Proprietary" with "Pty", for example. And Centrelink assumed that its own data had been entered correctly, meaning that a simple typing mistake one fortnight would ruin the data-matching process.

"Anyone who designs systems that deal with data should be aware of the issue of data entry errors. If you're smart, you try to ensure the highest quality of data at the point of entry as you can, because every time you manipulate the data, you risk adding errors. If your data starts off messy, then it just gets messier later on," Warren wrote.

"Which is exactly what is happening with Centrelink's data matching."

Most of these problems can be put down to mere stupidity and incompetence. But two issues would seem to indicate an ethical vacuum.

- Centrelink failed to plan for the volume of challenges coming in. According to the website [#NotMyDebt](#), which has been tracking the debacle, “Centrelink will not allow you to speak to their in-office staff regarding the debt, and insist that you lodge your appeal online (the website’s reportedly frequently offline) or via the telephone (people report it is very hard to get through).”
- While Centrelink has allocated resources and priority to chasing down overpayments, no attempt has been made to identify underpayments.”

“Everyone agrees that the government has an obligation to identify any overpayments. But what has changed is that the Turnbull government has removed the human oversight and let loose a poorly designed computer algorithm, effectively outsourcing the verification to the recipient,” [#NotMyDebt](#) writes.

This is precisely the issues discussed by Nigel Phair elsewhere in this issue of DirectorTech.



The BS at the core of big data, and its risk

“Given the rise of big data as a socio-technical phenomenon, we argue that it is necessary to critically interrogate its assumptions and biases,” wrote danah boyd and Kate Crawford in their 2012 paper [Critical questions for big data](#).

“Too often, big data enables the practice of apophenia: Seeing patterns where none actually exist, simply because enormous quantities of data can offer connections that radiate in all directions. In one notable example, Leinweber (2007) demonstrated that data mining techniques could show a strong but spurious correlation between the changes in the S&P 500 stock index and butter production.”

They derided what they called the core mythology of big data:

“The widespread belief that large data sets offer a higher form of intelligence and knowledge that can generate insights that were previously impossible, with the aura of truth, objectivity, and accuracy”.

One clear example is the so-called “talent analytics” used in recruitment. The idea is that if you analyse the behaviours of potential employees — whatever data you can get — it will tell you who’s the best for their needs.

The popular name for such statistical analysis right now is “machine learning”, and one of the problem with machine learning is that you can’t always know exactly how it’s reaching its conclusions.

Has the algorithm included data on race, sex, disability, age, or marital status, putting you in breach of anti-discrimination laws, for example?

Big data ethics is one for the board

Clearly, there are risks involved when an organisation appoints computers as managers, operating according to algorithms and a big pool of data.

In nearly every corporate conference on big data I've been to, speakers have acknowledged the ethical issues too. But what are they doing about it? I regularly ask them a key question.:

What organisational structures and processes do you have in place to ensure that those issues are on the table at every stage, if any?

Usually the answer implies that the organisation may well have a compliance committee that looks at corporate compliance committee which looks at regulatory issues such as privacy and consumer law.

But the answer rarely mentions any specific policy or process.

If your organisation is already rolling big data and algorithmic systems, then the time as come to include ethical considerations in the mix.

Disclosure: Stilgherrian travelled to Ho Chi Minh City as a guest of the Asia-Pacific Network Information Centre (APNIC), whose conference was held in conjunction with APRICOT.



Additional Resources

- Full video: [Tim O'Reilly's keynote address to the APRICOT/APNIC conference.](#)

001.008

Plan for your inevitable cybersecurity breach

*You'll recover faster if you've prepared a cybersecurity incident response plan, writes **Nigel Phair**. Maybe you even need a CERT.*



It's called incident incident response (IR) management in the cybersecurity world: detecting and responding to computer and information security events, and protecting organisational assets and systems from further interference.

Global computer networks, mobile phone services, and e-commerce have created an astonishing interconnectedness. IT systems and networks were already the backbone of most businesses, providing communications, storage, business functions, and administration. All these new connections are a boon for both business and the end user, but they also bring system vulnerabilities.

As a result, the number of cybersecurity incidents effecting businesses is on the rise.

According to Australia's **2015 Cyber Security Survey: Major Australian Businesses** (PDF), some 50 percent of businesses surveyed experienced at least one cyber incident in the previous 12 months, up from 39 percent the previous year. As the SANS Institute's **Incident Handler's Handbook** (PDF) therefore reminds us, it's "a matter of when, not if, a compromise or violation of an organisation's security will happen".

The constant stream of new cyber threats means that incident response often feels like a brand new encounter every day. An effective IR capability must therefore respond rapidly and effectively to those threats.

This challenge is usually addressed by setting up a CERT, a Computer Emergency Response Team, sometimes called a Computer Emergency Readiness Team.

The US Department of Homeland Security's US-CERT **defines** a CERT as "a concrete organisational entity (i.e. one or more staff) that is assigned the responsibility for coordinating and supporting the response to a computer security event or incident."


Any organisation can use the CERT concept, including governments, commercial businesses, and educational institutions.

At a minimum, a CERT should respond to incidents as they arise — but even that simple-sounding task can include numerous responsibilities as part of the organisational response.

- **Legal and Policy:** What are the organisation's obligations? Have there been breaches of the *Privacy Act*? Who needs to be informed? What are the policy implications, particularly if the attack is against a government entity?
- **Communications:** What is the messaging to the media, staff, customers and stakeholders? Who is the spokesperson, and what are they authorised to say?
- **Technical:** What is the in-house capability to respond technically? Do outsourced providers need to be brought in?

If it's going to perform effectively, the CERT team needs a solid foundation, so the first phase of IR is creating that foundation. Setting up the CERT should be a process based on communication, planning, and even rehearsal of the IR procedures themselves.

Once the organisation becomes aware of the breach of a critical system, the IR process moves through three further phases, all of which needs to be planned for: analysing the breach once it has been detected; deciding upon and implementing an incident containment and eradication strategy; and finally, feeding the lessons learned from that breach back into the IR plan, and the organisation more generally.



Developing and drafting an incident response plan is critical to ensuring the organisation can at least start the response.

An effective initial response is important, and will help in the later phases. Without understanding the incident, it may be difficult to limit the damage or prevent any further damage from happening.

The post-incident activities phase — putting practices in place to prevent, or mitigate, that particular kind of incident in the future — will also benefit from having documented the incident.

The Australian 2015 Cyber Security Survey notes that 77 percent of responding businesses have cyber security incident response plans in place, with 82 percent using externally-developed IT security standards or frameworks in preparing that plan.

The incident response plan can be further enhanced by taking some additional steps in its development.

- Know your data. It helps to know what data is exposed by your systems and where, and it helps to know what data might have been targeted in a particular breach.
- Have a strategy in place for contacting stakeholders or entities affected by the security incident.
- Make sure your organisational CERT understands threat management, both before and after an incident occurs.



001.A1

About the Contributors

Meet this issue's writers.



Justin Warren @jpwarren

Justin is a 20-year veteran of the IT industry with a background in infrastructure and software development. He's a consultant and analyst to companies large and small, as the owner of Melbourne-based PivotNine Pty Ltd. He also writes about the business of IT, with particular interests in infrastructure, security, open source software, and marketing.

www.eigenmagic.com



Kate Carruthers @kcarruthers

Kate is Chief Data Officer for UNSW Sydney, and an Adjunct Senior Lecturer in the School of Computer Science & Engineering at UNSW Sydney. She is a member of the Telstra Industry Advisory Board and the NSW Government's Data Analytics Centre Advisory Board. She has extensive experience in senior executive roles for diverse organisations including Citibank, General Electric, AMP, Westfield (now Scentre Group), and NSW Treasury. katecarruthers.com



Stilgherrian @stilgherrian

Stilgherrian is a freelance journalist, commentator and broadcaster who covers the politics of the internet, cybersecurity, cybercrime, and privacy. He studied computing science and linguistics before a wide range of media and technology work. He currently files for ZDNet, Crikey and others, is a frequent guest on radio and TV, and was technology consultant for the recent ABC TV drama *The Code*. stilgherrian.com

001.A2

About DirectorTech

Who we are, and how to find us.

DirectorTech is a quarterly advisory aimed at company directors of government entities, listed companies, private organisations and not-for-profits. The content for each advisory is researched, analysed and written by a diverse range of international contributors, each expert in their field. Advisories include written content specifically tailored for company directors and individual interactive webinars which focus on identified organisational issues.



Nigel Phair, Managing Editor and Publisher @nphair

Nigel is an Adjunct Professor at the Centre for Internet Safety at the University of Canberra, and is an influential analyst on the intersection of technology, crime and society. In a 21 year career with the Australian Federal Police he achieved the rank of Detective Superintendent and headed up investigations at the Australian High Tech Crime Centre for four years.



DirectorTech is an initiative of the **Centre for Internet Safety** at the University of Canberra.

Web: directortech.com.au

Email: admin@directortech.com.au

Phone: +61 408 437 056

Twitter: [@director_tech](https://twitter.com/director_tech)

Centre for Internet Safety
University of Canberra ACT 2617
Australia

ABN 28 150 165 343

DirectorTech is ©2017 Centre for Internet Safety. All rights reserved. Copyright in individual contributions is retained by the authors.

Cover and internal digital illustrations ©2017 Stilgherrian. All rights reserved. Image for 001.001 includes **ONR explains Project BlueShark** from the Wikimedia Commons, in the public domain. Image for 001.003 includes a **map** from Uber. Image for 001.008 includes an US Air Force image of personnel of the 624th Operations Center, located at Joint Base San Antonio - Lackland, in the public domain.