

Data sovereignty in action: Designing, building and implementing a radically distributed health information system in Aotearoa New Zealand

ALEX POOR

A thesis submitted to Auckland University of Technology in fulfilment of the requirements for
the degree of Doctor of Philosophy (PhD)

2022

School of Social Sciences and Public Policy

Primary Supervisor: Professor Marilyn Waring

Secondary Supervisor: Professor Rhema Vaithianathan

Third Supervisor: Professor Dave Parry

Acknowledgements

Firstly I acknowledge and thank my supervisors – Professor Marilyn Waring, Professor Rhema Vaithianathan and Professor Dave Parry – for their patience, support and guidance. Each of you experts in very different fields, you were vital in helping me to finally craft the thesis that I always wanted to write; one that I hope sits comfortably across disciplines, rather than falling between them.

Secondly, a big thanks to the Holochain community; a warm and welcoming space. David Atkinson, Eric Bear and Art Brock particularly went out of their way to talk and connect me with others. Bear really pushed to help me out when I quickly got stuck with Holochain, and found my Holochain mentor and sensei – Nastasia Emelianova – who volunteered to spend many many hours on Zoom dragging me through learning Rust and architecting distributed applications. Nastasia – this thesis would have been absolutely impossible without you, and I just don't know how to thank you for your generosity and patience.

Thirdly, I sincerely thank all the interview participants who also gave their time freely to contribute to this research. There was a real breadth of expertise, and sitting in the middle of it all was quite overwhelming but an amazing learning experience. Thank you all for the robust conversation.

Finally, to Jessica, it must have felt like it would never end... Thank you for believing in me.

Contents

List of Figures	vii
List of Tables	ix
Acronyms	xi
Glossary	xiii
Declaration	xvi
Abstract	xvii
1. Introduction	1
1.1. Rationale and significance	2
1.2. Research methodology	4
1.3. Research questions	5
1.4. Research aims and objectives	6
1.4.1. Understanding public attitudes	7
1.4.2. Demonstrating technical feasibility	8
1.4.3. Identifying implementation opportunities and challenges	9
1.5. Thesis structure	10
2. Literature review	12
2.1. Key concepts	12
2.1.1. Data sovereignty	12
2.1.2. Data privacy	15
2.1.3. Centralisation	17
2.1.4. Decentralisation	19
2.1.5. Distribution	21
2.2. Perspectives on data sovereignty	22
2.2.1. Geopolitical data sovereignty	22

Contents

2.2.2.	Indigenous and Māori data sovereignty	27
2.2.3.	Individual perspectives	31
2.2.4.	International examples	44
2.2.5.	The Aotearoa New Zealand context	56
2.3.	Centralisation, decentralisation and distribution	59
2.3.1.	Understanding network types	59
2.3.2.	Why is distribution important?	67
2.3.3.	Candidate distributed approaches	72
3.	Methodology and Methods	88
3.1.	Design Science Research (DSR)	88
3.1.1.	Artifacts	89
3.1.2.	Theory vs action	90
3.1.3.	The DSR process	91
3.2.	Research approach	94
3.2.1.	Problem provenance, rationale and significance	94
3.2.2.	Feasibility	94
3.2.3.	Scope	95
3.2.4.	The DSR threshold	95
3.2.5.	Define requirements	95
3.2.6.	Define possible/alternative solutions	96
3.2.7.	Explore knowledge base support of alternatives	97
3.2.8.	Prepare for design/evaluation	97
3.3.	Research phases	98
3.4.	Ethics Approval	98
3.4.1.	Consumer Panel Survey	98
3.4.2.	Interviews	100
4.	Research Phase 1: Consumer Panel Survey	101
4.1.	Goals and objectives	102
4.2.	Definition of key concepts	103
4.3.	Generation of hypotheses	103
4.4.	Choice of survey mode	103
4.5.	Question construction	104
4.6.	Sampling	107
4.7.	Administration and data collection	110

Contents

4.8.	Summarisation and analysis	111
4.8.1.	Quantitative analysis of structured question responses	112
4.8.2.	Qualitative analysis of unstructured question responses	118
4.9.	Conclusions and communication of results	122
5.	Research Phases 2-3: V1 Prototype design, build and evaluation	124
5.1.	Defining requirements	124
5.1.1.	Controlling access to data	125
5.1.2.	Auditing data access	125
5.1.3.	Correcting data	126
5.1.4.	Security	127
5.2.	Prototype requirements	127
5.3.	V1 Prototype design, build and evaluation by requirement category	129
5.3.1.	Recording and retrieval (RP1-3) – Data	130
5.3.2.	Recording and retrieval (RP1-3) – Functions	132
5.3.3.	Recording and retrieval (RP1-3) – Testing / evaluation	134
5.3.4.	Access control (RP4-6) – Data	140
5.3.5.	Access control (RP4-6) – Functions	145
5.3.6.	Access control (RP4-6) – Testing / evaluation	146
5.3.7.	Audit (RP7) – Data	149
5.3.8.	Audit (RP7) – Functions	150
5.3.9.	Audit (RP7) – Testing / evaluation	151
5.4.	Summary	153
6.	Research Phases 2-3: V2 Prototype design, build and evaluation	155
6.1.	V2 Prototype requirements	156
6.2.	User interface design, build and evaluation (RP8)	157
6.2.1.	Choice of framework	157
6.2.2.	Frontend design	158
6.2.3.	Ex ante evaluation	161
6.2.4.	Frontend development and ex post evaluation	164
6.3.	Summary	170
7.	Research Phase 4: Interviews	172
7.1.	Choice of method	172
7.2.	Participants	174
7.2.1.	Participant selection criteria	174

Contents

7.2.2.	Sample size	178
7.2.3.	Participant engagement	178
7.3.	Interviews: Design and implementation	182
7.3.1.	Developing the interview protocol	183
7.3.2.	Participant interviews	188
7.4.	Ethics	188
7.5.	Transcription	189
7.6.	Analysis	189
7.6.1.	Understanding the data	191
7.6.2.	Initial coding	193
7.6.3.	Search for themes	195
7.6.4.	Review themes	198
7.6.5.	Define and name themes	208
7.7.	Discussion	211
7.7.1.	Access to data	211
7.7.2.	Capability	218
7.7.3.	Control over data	224
7.7.4.	Trust	236
7.8.	Summary	243
8.	Discussion	245
8.1.	Research questions	246
8.1.1.	How do people feel about the ability to have absolute control over their personal data?	246
8.1.2.	How can technology support the distribution of personal health data and a move toward data sovereignty?	248
8.1.3.	How could a distributed health information system be implemented in Aotearoa New Zealand?	249
8.2.	Implementation model	253
8.3.	Significance of research	253
8.3.1.	Contribution to the knowledge base	256
8.3.2.	A call to action	257
8.3.3.	Possibilities for further research	257
8.3.4.	Conclusion	259
	References	260

Contents

A. Survey Participant Information Sheet	287
B. Survey: EA1 Approval	291
C. Survey final report	294
D. Interviews: EA1 Approval	303
E. Interview Participant Information Sheet	305
F. Initial coding from Research Phase three	309

List of Figures

2.1. Sample data flow across healthcare entities	19
2.2. X-Tee, the technical layer of E-Estonia (Republic of Estonia Information System Authority, 2022).	45
2.3. Three elements of the Data Embassy Initiative. Sourced from Estonia Ministry of Economic Affairs and Communications (2016).	47
2.4. vTaiwan – how it works	55
2.5. Centralised, decentralised and distributed networks (Baran, 1964)	60
2.6. Health sector layers, adapted from Baran (1964).	61
2.7. Blockchain as a decentralised network	64
2.8. Git workflow	69
2.9. BitTorrent as a supervised distributed system	72
2.10. Example Hyperledger Fabric blockchain network	75
2.11. Holochain Distributed Hash Table	80
3.1. Research phases and outputs	99
4.1. Ethnicity comparison of Consumer Panel and survey respondents	112
4.2. Gender comparison of Consumer Panel and survey respondents	113
4.3. Age group comparison of Consumer Panel and survey respondents	113
5.1. Tryorama test result for RP1-3	141
5.2. Holochain Distributed Hash Table	142
6.1. Wireframe: Login	159
6.2. Wireframe: Main/data	160
6.3. Wireframe: Access	161
6.4. Wireframe: Audit	162
7.1. Assessment of participants against selection criteria	181
7.2. Initial thematic map	197
7.3. Thematic map: revision one	202

List of Figures

7.4. Thematic map: revision two	207
7.5. Final thematic map	210
7.6. Consent and social licence conflation under centralisation	227
7.7. Separation of consent and social licence under distribution	227
7.8. Automation (xkcd, n.d.)	237
8.1. RDHIS implementation model	254
8.2. Analytical communities	255

List of Tables

2.1. Types of data sovereignty. Adapted from https://www.temanararaunga.maori.nz/ (accessed 20 April 2022).	14
2.2. Public attitudes to sharing data – key themes, adapted from Aitken, de Jorre, Pagliari, Jepson, and Cunningham-Burley (2016).	32
2.3. Anonymisation problems with OSN data	42
2.4. Centralisation matrix	66
2.5. Comparison of blockchain types. Adapted from H. Wang, Zheng, Xie, Dai, and Chen (2018).	76
2.6. Blockchain challenges in healthcare. Adapted from de Aguiar, Faical, Krishnamachari, and Ueyama (2020).	78
2.7. Comparison of candidate approaches	86
3.1. Artifact forms, as described in March and Smith (1995)	90
3.2. Design Science Research Process. Adapted from Peffers et al. (2006).	92
3.3. Design Science Research Roadmap. Adapted from Gazem, Rahman, Saeed, and Iahad (2018).	93
3.4. Study requirements. Based on Gazem et al. (2018).	96
4.1. Stages of survey design. Adapted from Lavrakas (2008).	102
4.2. Survey question standards. Adapted from Groves et al. (2009).	105
4.3. Conventional pretesting results	105
4.4. Key components of effective question construction. Adapted from Cowles and Nelson (2015).	106
4.5. Consumer Panel gender	109
4.6. Consumer Panel age group	109
4.7. Consumer Panel ethnicity	109
4.8. Top five respondent subgroups	114
4.9. Responses to survey – section two	114
4.10. Responses to survey – section three	116

List of Tables

4.11. Top ten key concepts from First Cycle coding	119
4.12. Final coding concepts used	120
5.1. V1 Prototype requirements	128
5.2. Design Evaluation Methods, adapted from Hevner, March, Park, and Ram (2004).	136
6.1. V2 Prototype requirements	156
6.2. Ex ante evaluation findings	163
6.3. Testing dimensions for web apps, adapted from Kinsbruner and Bahmutov (2022).165	
6.4. V2 group testing session requirements.	169
7.1. Health data stakeholders. Adapted from The European Institute for Innovation through Health Data (2021).	175
7.2. Initial participant list and selection criteria	180
7.3. Interview question types. Adapted from Rubin and Rubin (2012).	185
7.4. Reflexive Thematic Analysis process. Adapted from Braun and Clarke (2006). .	191
7.5. Initial themes from interview transcripts	192
7.6. Thematic map: revision one	200
7.7. Testing themes are clearly defined	209
7.8. Stages of decentralisation	216
7.9. Key findings from Research Phase 4	243
F.1. Thematic Analysis: Initial coding	309

Acronyms

API Application Programming Interface

CCP Chinese Communist Party

CS Computer Science

CSO Civil Society Organisation

DHT Distributed Hash Table

DSR Design Science Research

EHR Electronic Health Record

FHIR Fast Healthcare Interoperability Resource

FOSS Free and open source software

GDPR General Data Protection Regulation

GP General Practitioner

IDI Integrated Data Infrastructure

IPFS Interplanetary File System

IS Information Science

MELAA Middle Eastern, Latin American or African

MHN Midlands Health Network

MOH Ministry of Health

NFC Near-Field Communication

NFT Non-Fungible Token

Acronyms

NGO	Non-Governmental Organisation
NHI	National Health Index
NMDS	National Minimum Data Set
OSN	Online Social Network
PHO	Primary Health Organisation
PKI	Public key infrastructure
PMS	Patient Management System
PoW	Proof of Work
RDHIS	Radically Distributed Health Information System
RTA	Reflexive Thematic Analysis
SME	Subject Matter Expert
SSB	Scalable Secure Scuttlebutt

Glossary

API A service which permits interactions between computing systems or services. For example, a developer building a dashboard for a pizza delivery firm may access the Google Maps API to programmatically show a real-time map of current orders, or to display shop locations.

Centralisation A paradigm where ICT infrastructure and data is under the control of an organisation, removed from those who are generating the data.

Cloud Any ICT infrastructure or service which is delivered from an external data centre over the internet. This encompasses a very wide range of products and services, from server/-compute instances, to serverless code execution or block/object storage.

Consortium blockchain Similar to a Permissioned blockchain except only a small number of validating nodes exist, usually in a high trust environment.

CSO Analogous to an NGO (see below).

Decentralisation A paradigm where previously centralised data has been transitioned into the control of individuals, groups and communities who exert sovereignty and decision-making authority over those assets.

DHT A key-value pair store utilised in distributed systems.

DSR A research methodology which prizes the generation of novel artifacts to solve real problems in innovative ways.

EHR The systematised collection of patient health information, recorded and stored in a digital format.

FHIR A modern health data information exchange protocol, utilising a common API format.

Git Software used by development teams to track changes in code, and to keep an audit log of all actions taken. Repositories can be modified locally and merged safely back into the whole, thus providing an important decentralised workflow.

Glossary

- Gossip** A peer-to-peer communication concept, whereby peers in a network 'gossip' with each other in order to replicate, share or reconcile information.
- Hapū** A Maori social unit which may be considered analogous to a "sub-tribe" in English.
- IDI** A centralised government mega-database, operated by Stats NZ Tatauranga Aotearoa, which joins up all available administrative datasets and can be analysed at the level of an individual to track use of services over time. It is accessed by researchers and government agencies.
- IPFS** A peer-to-peer network for storing and sharing data in a distributed file system.
- IS** An academic field concerned with all aspects of information and data.
- Iwi** A Maori social unit, often translated into English as "tribe".
- New Public Management** A public sector model rooted in neoliberalism, which gained popularity during the 1980s and attempts to improve efficiency by importing private sector methods.
- NFC** A communication protocol for devices in very close proximity.
- NFT** A digital asset that can be bought and sold using cryptocurrency.
- NGO** Typically non-profit entities engaged in humanitarian or social sector work.
- NHI** A unique identifier, issued at birth, for all health system users in Aotearoa New Zealand.
- OSN** An umbrella term used to refer to social media networks and online communities.
- Permissioned blockchain** A blockchain managed by an organisation, where permission is required to join and transactions and capabilities are controlled.
- Permissionless blockchain** A blockchain where anyone can participate and view transactions (for example, Bitcoin and Ethereum).
- PHO** A statutory entity in the New Zealand health sector who provide management and shared services functions to general practice.
- PMS** A computer system generally used for administrative purposes in healthcare. For example, managing appointments and recording clinical data.
- Private cloud** ICT infrastructure services operated for a single organisation, usually from a data centre which might be located within Aotearoa New Zealand, or off shore.

Glossary

PoW A blockchain consensus protocol that requires nodes to compete in solving a cryptography puzzle so as to validate transactions. Some blockchains offer a small reward to the winner (in Bitcoin this is known as "mining").

Public cloud Services rendered over a network that is open for public use. Examples being Amazon Web Services or Google Cloud Platform.

RDHIS A health information ecosystem where centralised silos are removed and data is owned and controlled by individuals.

RTA A qualitative text analysis methodology which aims to be transparent about bias.

SSB A distributed gossip protocol for peer-to-peer communication.

Web3 An iteration of internet technology, purportedly based on concepts such as decentralisation and distribution, but popularly focused on blockchain and cryptocurrency.

Declaration

“I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.”

Signed: 

Alex Poor

Date: 25th May, 2022

Abstract

This thesis examines the nascent, but rapidly maturing, ‘Web3’ technology and seeks to understand the role it can play in realising data sovereignty in the context of the New Zealand health sector. The data sovereignty movement has generated its own momentum, particularly within the indigenous literature, but it is still missing a *technical* means by which data can be brought under the control and ownership of groups and communities. I argue in this thesis that one such Web3 technology – Holochain – offers this potential.

The research has three distinct phases. First, a health consumer panel survey was completed which demonstrated support for the ownership and control of personal health information and some dissatisfaction with the status quo. Respondents were also asked to imagine a data sovereignty-focused health information system, and what that should look like. These insights became user requirements, which were utilised in the second research phase to build a prototype Holochain app. Formal evaluation found it met all stated requirements for a fully distributed health information app. Having established that data sovereignty can be meaningfully supported by this new technology, I turn to look at what a radically distributed health information system would look like from a policy perspective.

This third phase was completed by interviewing experts from a broad range of disciplines. There are certainly many obstacles to implementation of data sovereignty utilising distributed technologies, but these primarily exist in the unravelling of the status quo. I frame this status quo as being entirely dominated by the extant model of data management, which relies on centralising data into silos – the *centralised hegemony*. There are also significant issues to overcome around maintaining essential government access to data, but I argue that this is merely a trust and social licence issue which government should in any case be focused upon.

Data sovereignty represents a paradigm shift in management and use of data. I contend that decentralised and distributed models have not been part of the discussion before now, because of the relative immaturity of these technologies. Having demonstrated in this research that it *can* work, however, I propose that the default to centralisation is urgently reconsidered. This research fills an important gap in the literature by demonstrating that a feasible technical solution for data sovereignty exists, and by specifying the policy steps needed to unravel the centralised hegemony.

1. Introduction

“Miserliness leads to great expense, hoarding leads to deep loss” (Dao de jing, 44).

This thesis investigates the concept that personal health data can be owned and controlled by individuals, groups and communities rather than by government agencies or for-profit corporations. This concept is very much aligned with the varying strands of ‘data sovereignty’, a concept that is gaining momentum within New Zealand society predominantly from a te ao Māori perspective.

The prevalent model of general data management today relies exclusively on *centralisation* – a pattern that sees personal data aggregated into silos, over which an individual has relatively little oversight and almost no control. The discourse then shifts to interoperability (*how can we connect these data stores?*), yet individuals will still exert only minimal control over their data. There are good reasons for this. There is no extant technology available which could support a different data management pattern right now. But such nascent technologies do exist, and they are maturing rapidly.

In this thesis I will argue that the centralisation pattern has become so embedded that it constrains our thinking around personal health data sovereignty. If data sovereignty can be defined as “holding decision-making authority over one’s data” then, within a centralised paradigm, our solutions will be limited to forming advisory groups, increasing governance oversight, conducting consultation and perhaps even giving users an idea of how their data is being shared and used. But it clearly does not get anywhere close to a notion of real *sovereignty*¹. My thesis is that sovereignty can only be achieved when personal health data resides solely with the owner (the person it concerns), or the owner’s group, community, iwi or hapū – however they define it to be meaningful and relevant. That is, there would be *no more centralised data stores* at all.

This will undoubtedly sound far-fetched, and not merely because centralisation has become so embedded. It is true that there is a debate to be had about who really does *own* health data – *isn’t it at least co-created?* But I would value this opportunity to at least ask some provocative

¹Please note that my working definition and representation of data sovereignty differs from the literature that does exist. For example in Kukutai and Taylor (2016) the focus is on sovereignty as a governance issue. My representation here of “real sovereignty” is not intended to challenge that definition. This will be discussed in more depth in chapter 2.

1. Introduction

questions and, as a central component of this research, to demonstrate that there are already tools which are mature enough for us to start planning seriously for a distributed health information future. While the focus within this thesis is on the health sector, it will be obvious that there are similar implications for personal data in any context and I anticipate the findings will be somewhat generalisable.

At the culmination of the thesis, I will have proven that distributed health data is technically possible, that it can enable data sovereignty and that both align with public sentiment. Together with a summary of challenges and opportunities in implementation (please see chapter 7 on page 172), I will present the feasibility and applicability of what I will call a ‘Radically Distributed Health Information System’ (RDHIS) for Aotearoa New Zealand.

1.1. Rationale and significance

At the outset of my PhD journey, my original area of research interest was to investigate the current state of ‘data sovereignty’ in the Community Provider/Non-Governmental Organisation (NGO) sector of Aotearoa New Zealand. I was interested to find out how much control individuals actually have over their own data, particularly in a sector where people often seek support for very personal and sensitive issues. This was a topic of great interest to me personally, since I have for a long time been involved in the free and open source software community and have also worked extensively with data in the health sector. I should also add that I am a person living with a chronic disease, and have a lived experience of being unable to access and use personal health data in a way that would be meaningful to me.

It is important to note here also, as early as possible, that the term ‘data sovereignty’ is a rather slippery concept and means different things to different groups. I will reference the vital contributions of indigenous and Māori data sovereignty researchers but, as a Pākehā, this is not my focus. I am interested in data sovereignty as it might apply to any individual or group in society as best meets their needs. It is *having the ability to control your personal data the way you want to*. I will discuss further some definitions and key concepts in section §2.1.

My initial research and literature review led me to the conclusion that there is no real data sovereignty anywhere in Aotearoa New Zealand². This is not necessarily a surprise; the concept is still relatively new. But there are many *related* concepts that have been floating around our digital roadmaps and statements of intent for a very long time – patient-centred care, single electronic health records, empowering patients with data, *et cetera*.

The Ministry of Health (MOH) have published a range of strategic reviews, IT plans, roadmaps and digital portfolios since 2001 – all of them noting both the need to empower ‘consumers’ and

²Although there is certainly a lot of impressive work happening, which may be considered a step towards it.

1. Introduction

the importance of data³. The 2016 New Zealand Health Strategy specifically notes the importance of being patient-centric, and that we should be “making New Zealanders ‘health smart’; that is, they can get and understand the information they need to manage their care” (Ministry of Health, 2016, p. 16).

This takes shape in the MOH *Hira* programme, which is a digital transformation project aimed at enabling a “whole new way for different data systems to connect ... [which would] ... empower people and their whānau to better manage their health, wellbeing and independence” (Ministry of Health, 2021b, p.6). The business case specifically mentions data sovereignty several times but, interestingly, only in the context of Māori and only as a non-technical component. This is not at all a criticism of such an ambitious programme of work, but it does set the scene for how the possibilities for data sovereignty are being constrained by the centralised hegemony.

At the same time, the mHealth market is growing and it is not unusual for a person to use several health-related devices or apps. A user with a FitBit device, for example, may also use a so-called ‘femtech’ app to monitor menstruation or ovulation, with the data from each being ingested into an aggregator service such as Apple Health⁴. Already, this one person’s health data resides in three separate commercial vendor systems and this is not to mention some alarming privacy issues that are thrown into the bargain (Rosato, 2020; Schiffer, 2021).

So there is an additional concern that the ability to hold and share personal health data (but not control it) has been captured by the private sector, and even powerful opportunities to share patient-generated data with health professionals relies on connecting to the centralised data stores of app vendors or aggregators who are either profit-driven, or whose primary business is advertising. Given this, it seems that there is a place for a public health sector response to the gap that these commercial products are filling. *Hira* offers a roadmap towards interoperability and connectivity, but does not solve the sovereignty issue around this most sensitive of data.

Internationally, the most innovative initiatives all focus on management and sharing of data only *once it has already been collected*, reinforcing centralisation (for example, Estonia’s ‘X-Tee’ system). The end user, or patient, still has no voice in terms of critical concepts such as:

- Deciding how, or whether, their information is shared, and
- Deciding how it is utilised or analysed, so that marginalised groups are not misrepresented, under-represented or somehow marginalised further.

Achieving these two things is not possible currently from a technical perspective, and so the focus has been on consensus building, changing governance arrangements and defining optional

³This background will be discussed further in chapter 2.

⁴Similarly Facebook has intermittently made a push into the healthcare vertical (Reader, 2021), which it can use to complement more than 1250 petabytes (or 1.22 exabytes) of personal data (Pan et al., 2021) already being used for advertising profit.

1. Introduction

standards or operating principles. These are all very laudable initiatives, but they do not represent real data sovereignty. Decision making power cannot be reliably exerted when someone else is holding your data.

A technical shift is occurring, though, with the vaunted ‘Web3’ paradigm which, although popularly focused on cryptocurrency and blockchain, encompasses a wide range of new technologies that do offer hope to data sovereignty advocates. There is very little research or literature on these, outside of blockchain, and none in a healthcare setting. There is definitely a ‘gap’ in the literature, therefore, and I hope to show what technical frameworks offering true distribution of data could achieve in health.

There are, of course, many issues with distributing data purely because we have operated under a centralised paradigm for such a long time. I will review these issues and, finally, present a way forward for distribution of data in the Aotearoa New Zealand health sector as a means to achieve data sovereignty.

1.2. Research methodology

In the previous section I identified a *technical* gap, where I described how I had come to the realisation that nascent distributed technologies can be practically applied to meaningfully advance the data sovereignty debate. This will be discussed extensively in chapter 2, but I must state here that what I felt most passionate about doing as part of this research was to *build something*. This has important ramifications for the process of selecting methodology and the thesis structure.

Right at this early stage, it was clear to me that Design Science Research (DSR) would be utilised as the main methodology. DSR attempts to bring academic and methodological rigour to a design and development process (predominantly within Computer Science disciplines) that has been criticised as being “mostly driven by intuition ... leading to a rather inefficient and tinkering-based process” (Schork & Kirchner, 2018, p.2). Information Systems (IS) research can often be descriptive and informal, so DSR’s focus on *utility* helps to form a bridge between IS research and practice (Peppers et al., 2006).

The focus of DSR, however, remains the production of an *artifact*. It is an approach “that invents, or builds new, innovative artifacts for solving problems or achieving improvements” (Iivari & Venable, 2009, p. 3) and thus lends itself well to digital products. I will utilise a variant of DSR proposed by Gazem et al. (2018), who set out 14 discrete steps to complete prototype development with methodological rigour.

According to Hevner et al. (2004), what sets DSR apart from a “tinkering-based process” is formal evaluation. That is, clear functional requirements must be generated from primary research and these must form the basis of prototype evaluation. These requirements will be

1. Introduction

generated by the consumer survey discussed in section 1.4.1. The findings from both quantitative and qualitative analysis of those results will be developed into a requirements framework, against which the prototype will be evaluated. The evaluation process itself takes several different forms as I move through the DSR process. This is discussed in more depth in chapter 3.

Finally, I will attempt to locate distribution of data in a real world context by seeking the thoughts of experts. For this research phase I will utilise Reflexive Thematic Analysis (RTA). This was chosen because of the large amount of qualitative data that this phase generated, and my concern about how subjective the coding process felt. Specifically, what if I am unconsciously ignoring a theme that doesn't fit my own world view? Or, what if I attribute disproportionate weight to themes that support my personal beliefs? RTA helps the researcher deal with this by asserting that there is no such thing as a positivist truth that is simply waiting to be discovered (Braun & Clarke, 2006). Understanding is inherently influenced by our own thoughts and biases – we simply need to name these, understand them, and explain to the reader why we have come to a certain decision.

I should note that the choice of methodology here might encourage us to overlook entirely any examination of epistemological or ontological features of the research subject, which would ordinarily precede selection of methodology and then inform methods. In discussing the need for RTA, in the previous paragraph, I have quite clearly stated my personal alignment with constructivism by seeking out a method which rejects the idea of positivist truth. Conversely the analysis of survey data, discussed in section §4.8, utilised a mixed methods approach which saw quantitative analysis followed up with content analysis of free text responses. DSR is more or less agnostic about one's epistemological position, and merely asks that a unique problem can be solved by production of an artifact whilst following a structured process.

1.3. Research questions

A clear research question is the essential starting point for development of aims, objectives and, thus, a logical and consistent thesis (Martindale & Taylor, 2014). In chapter 3 I will discuss use of the DSR methodology, and so it is interesting also to note the assertion of Chatterjee and Hevner (2010) that defining the research question is the most important part of the DSR process.

Three basic functions of research questions particularly apply to DSR projects:

- Defining scope
- Driving the research process, and
- Positioning a project's contributions (Thuan, Drechsler, & Antunes, 2019).

1. Introduction

However DSR also offers two additional functions by:

- Ensuring that innovative artifacts are produced by posing innovative research questions⁵, and
- Addressing a knowledge gap by asking relevant research questions (Gregor & Hevner, 2013).

I certainly hesitate to declare that my research questions are innovative, but I am confident at least that this thesis fully meets the DSR requirement that it “addresses important unsolved problems in unique or innovative ways” (Hevner et al., 2004, p. 81). There are three research questions:

1. How do people feel about the ability to have absolute control over their personal data?
2. How can technology support the distribution of personal health data and a move toward data sovereignty?
3. How could a distributed health information system be implemented in Aotearoa New Zealand?

Collectively these research questions have defined the scope (a focus on health data in Aotearoa New Zealand) and confirmed that the feasibility of data distribution will be tested and placed in a real world context to consider implementation. Finally, the aims and objectives (see below) for research question two make clear that I am proposing to develop an innovative artifact and that an interesting new contribution to the knowledge base will be made.

1.4. Research aims and objectives

In this thesis I will demonstrate that distribution of personal health data is technically feasible and that it should start to inform our discourse around data sovereignty.

As I have mentioned, data sovereignty can have many meanings and there are certainly many threads to this topic. I have identified three main aims that link to the research questions that this thesis will address:

1. Understanding public attitudes toward data sovereignty, and identifying demand or levels of support

⁵DSR by definition eschews anything that is not an original contribution to the knowledge base, or solving a new problem in an interesting way.

1. Introduction

2. Demonstrating that distribution is technically feasible, and can underpin a practical move toward data sovereignty
3. Identifying how distribution of data might be implemented in the Aotearoa New Zealand health sector.

I will next discuss each aim in more detail, together with how I propose to achieve them (the objectives).

1.4.1. Understanding public attitudes

It would be pointless to write a thesis about distribution of data if there was no evidence that it did anything that people wanted. This aim, therefore, is about establishing what public attitudes to data sovereignty are and whether or not there is indeed a desire for people to exert decision-making authority over their data. Put another way, is there *demand*?

Since data sovereignty is a relatively new concept – and centralisation has been so dominant – it is reasonable to assume that this topic is not prominent in the public consciousness. During this research I have not stopped ten people in the street and asked them what they think about data sovereignty, but I imagine that most people would not have seriously considered it. As I have mentioned, the Māori data sovereignty movement has generated excellent momentum and so my assumption here may be egregiously wrong when applied to a Māori audience. But this is exactly what I need to determine, both from the available literature and empirically; *who cares about this?*

There is a wealth of Aotearoa New Zealand and international literature to draw from, with many researchers identifying that people *do* want to exert control over their personal data (Aitken et al., 2016; Data Futures Partnership, 2017a; Garrison et al., 2016) whilst, at the same time, finding a generalised fatalism around the *status quo* (Malkin, Bernd, Johnson, Egelman, et al., 2018; McMullan, 2015; Solon, 2018).

A potentially confounding factor is that, as already mentioned, there is relatively little literature on data sovereignty (and personal data distribution, specifically). In reviewing the literature, therefore, an interpretive component is required to extrapolate from findings around more ubiquitous topics such as privacy or consent. On that basis, primary research is certainly required and, ideally, should be carried out within an Aotearoa New Zealand context. I will achieve this aim by meeting the following objectives:

- Conducting a literature review to ascertain the breadth of findings on topics related to data sovereignty

1. Introduction

- Conducting a survey of health ‘consumers’ to ask specific questions which relate to distribution as a way of realising data sovereignty.

1.4.2. Demonstrating technical feasibility

The centralised hegemony that I have already mentioned has enabled rapid and meaningful advances in technology over the last 50 years. When Edgar F. Codd began designing the origins of relational databases in 1970 (Codd, 1970) he could not, of course, have predicted the volumes of data that are being generated today. Ever since that time, the relational database has remained the centrepiece of any digital tool, application or software product.

Centralisation works *really well!* Developers and programmers are very familiar with its architectural approaches; managers and stakeholders value the security and stability that it appears to offer; network engineers are grateful to work with predictable flows of data. When you take that centrepiece away and let individual users connect directly with each other then, architecturally, what you have looks like absolute chaos. I am not saying distribution of data will be easy: it verifiably won’t.

But I argue that the benefits of centralisation will tend to accrue to those individuals working within or around that model – designers, developers, analysts, managers, government agencies. The stakeholders who are farthest away from these benefits will tend to be those people who have *generated* that data – citizens, individuals, regular people⁶. The point I am trying to make is that centralisation makes things easier for everyone *except* the individuals, or groups, who would like to make decisions about their personal health data and utilise it to increase wellbeing.

The concept of distributed data is so abstract (*Where does the data live? How do we access it?*) when viewed from within the centralised bubble that it appears very hard to know where to begin. Confusing our understanding of the concept further is what I will call its bastardisation, by increasingly popular technologies such as blockchain and cryptocurrency. Discussion of these is often accompanied by nebulous marketing terms such as “freedom” and “liberty” but, also very often, “decentralisation” and “distribution”. I will review this landscape in section §2.3, where I conclude that blockchain is ‘decentralised’ but not fully ‘distributed’.

But there are tools and technologies which offer a promising route to achieving distribution on a technical level. To demonstrate technical feasibility, I must work with an appropriate candidate to develop a proof of concept application that will perform all the functions we might reasonably expect. This aim will be achieved via the following objectives:

- Understanding functional requirements of a distributed app via primary research (consumer survey)

⁶It would be a step too far to suggest that individuals receive *no* benefits from centralisation.

1. Introduction

- Development of a prototype app using a distributed technology
- Evaluation of the prototype app.

1.4.3. Identifying implementation opportunities and challenges

Decentralisation is hard, certainly on a technical level. But if we are to move further towards tangible data sovereignty, it must somehow be situated within the real world.

I have already talked about the centralised hegemony, which I use to describe a dominant paradigm to which all surrounding processes and workflows have been bent. This is entirely natural, of course. If this is indeed “the way we do things”, then why wouldn’t our social policy, legislation, architectural models, business models and cultural norms reflect that? With this aim I am attempting to flesh out the rather obvious point that, if we want to achieve data sovereignty, then we need to implement distribution of data and, if we implement distribution of data, then lots of things will have to change.

Obstacles will be spread across a range of domains. For example, there will be many technical barriers to scaling such a concept nationally. The prototype app, being developed under ideal conditions, cannot fully consider the enormous fragmentation and complexity of interfacing with all the different public health information systems – many of which are still not ready to interoperate with external data sources, or each other, even today⁷.

Similarly there will be many political obstacles. A dominant theme throughout 13 in-depth semi-structured interviews with experts was a concern that distribution of data would cut off the centre’s access to administrative and other data that it currently relies upon⁸. Furthermore, the vendor market in healthcare can be extremely lucrative and vendors would be unlikely to open their systems to a service that by definition centres that data somewhere else. Research I carried out in 2016 highlighted the issue of ‘knowledge as a business model’ in the vendor landscape (Poor, 2016), but there is other literature agreeing that protecting revenue models represents a meaningful obstacle to technological change (Coiera, 2009; Debreceeny, Putterill, Tung, Gilbert, et al., 2002).

This aim therefore seeks to gather as much information as possible on what the challenges and opportunities of decentralisation are. Since there is a dearth of literature on this issue, I will conduct primary research and interview a panel of experts to seek their thoughts. The experts should naturally come from wide backgrounds and should have some stake in the topic of data sovereignty or decentralisation. To achieve this aim I will complete the following objectives:

⁷Interoperability and connectivity between health system data silos is a formal goal of *Hira*.

⁸This research phase is discussed fully in chapter 7.

1. Introduction

- Conduct in-depth semi-structured interviews with relevant experts, who can identify a range of opportunities and challenges with decentralisation
- Complete qualitative analysis on the outputs from this process, to identify key themes.

1.5. Thesis structure

My hope is that the structure of this thesis helps the reader to move alongside me on the research journey. Although there are four discrete research phases, each utilising differing methodologies and methods, I have attempted to keep them focused on the research questions and to order them in a logical fashion.

Chapter 1 (this chapter) aims to set the scene for the research. Why am I doing this, and is it actually something that needs investigating? I step the reader through the aims and objectives, and clarify the distinct research phases and why they are being carried out. I also briefly discuss the methodologies employed across all planned research phases.

Chapter 2 surveys the literature on data sovereignty. As already mentioned, data sovereignty specifically has a rather small research base thus far. Therefore, I look beyond this to related concepts such as consent and privacy to make inferences about data sovereignty. Since the surrounding concepts are potentially numerous, the literature review will focus on:

- Defining key concepts
- Individual perspectives and attitudes around data, data sharing and data use
- The current state of data management in the Aotearoa New Zealand health sector
- Innovative data management and data sharing models that have been implemented internationally
- Knowledge around the centralised and decentralised paradigms generally
- A review of candidate decentralised technologies.

Chapter 3 details the DSR methodology. I will review the methodology and explain its value and applicability to this research, before clarifying the key steps from the chosen model (Gazem et al., 2018). I will also clarify all research phases in this thesis, and how they relate to each other – attempting to tie them together into a coherent whole. Finally I will provide detail on ethical approval.

Chapter 4 describes the first research phase, where I survey a consumer panel to ascertain their thoughts and attitudes around data sovereignty. A basic content analysis approach is utilised for

1. Introduction

the qualitative component of this research phase, since the *corpus* was relatively small. This research phase has two main purposes:

1. To complement the literature review and develop understanding around significance of this topic
2. To generate functional requirements that can be utilised in the DSR process.

Chapter 5 discusses the first part of the DSR process. Functional requirements have been generated from the consumer survey, and these are a focus in terms of what the developed prototype must achieve. In this chapter I also spend some time discussing the development process, and explaining the inner design considerations for the decentralised app. This chapter aims to demonstrate the large amount of technical learning and work that went into that, as well as provide an insight into how decentralised app development might differ from the *status quo*. Finally, formal evaluation is carried out against the V1 prototype and learnings/findings are carried forward into the V2 design process (in the following chapter).

Chapter 6 summarises the evaluation results from the V1 prototype, and sets out what will be achieved in the next iteration of the app. V2 essentially entails development of a user interface for the prototype, which was not originally in scope but it was subsequently thought that omitting one would be an obstacle to learning, understanding and dissemination. Once again I go into some technical detail around how this was done, to reflect the large amount of learning I undertook to achieve it. The final evaluation for the V2 prototype takes the form of a recorded group test session, the video of which is available to view as a link in that chapter. A full link to the git repository containing all code is also given.

Chapter 7 discusses in detail the in-depth semi-structured interview process. I present the RTA methodology and explain its relevance, before stepping through the interview process. I attempt to take care during this chapter to specify all decisions and thought processes utilised during the coding process, so as to explain the outputs and theoretically make it more replicable. A wide range of themes were identified, and these are presented as a thematic map.

Finally, Chapter 8 draws all the research outputs together to try and ensure that the research questions have been addressed. I present the prototype app as a novel contribution to the knowledge base, being the first representation of a distributed health information app. But this must be located in a real world context, and I draw conclusions from the findings in Chapter 7 to do so.

2. Literature review

This literature review will build on the research questions outlined in section §1.3, and delve further into the issues surrounding each. As already noted, there is not a large amount of literature dealing specifically with personal data sovereignty, and the paradigms that may support it. However, a lot of research has been done on overlapping topics such as public attitudes to data sharing, privacy and consent.

This literature review plays a key role in answering the first research question – *How do people feel about the ability to have absolute control over their personal data?* – which will be complemented with primary research, to zoom in on the Aotearoa New Zealand health context. I will then switch focus to exploring the data management and data sharing landscapes in selected jurisdictions, before again bringing this back to the Aotearoa New Zealand context. Finally, I will explore current knowledge around the differing paradigms and bring this together to identify technologies that will be most suitable for use in developing a prototype artifact.

Firstly, however, I will clarify some key concepts. I have already used terms such as data sovereignty, and centralisation, but – since these are the foundation of this thesis – it is worth spending some time to further understand and define them.

2.1. Key concepts

2.1.1. Data sovereignty

There are two prominent schools of thought around the definition of data sovereignty, both of which are discussed further in section §2.2.

The first can be thought of as a *geopolitical* one, and has its roots in debate around national security. Prominent subjects of discussion in this paradigm relate to the following examples:

- Implementation of the US PATRIOT Act in 2001, following the 9/11 terrorist attacks
- The 2015 introduction of Federal Law 242-FZ in Russia, which mandates data localisation within Russian Federation territory
- The 2017 Cybersecurity Law in China, which also mandates data localisation within mainland China and provides for authorities to conduct checks on any company's operations.

2. Literature review

These are three examples of many¹ where legislation has been used to assert nation state sovereignty. In the Russian case, it was developed explicitly as a response to US sanctions and bilateral tension following Russia's illegal invasion and annexation of Crimea and the Donbas region of Ukraine in 2014 (Savelyev, 2016). In the case of China, the legislation specifically invoked the novel concept of '*cyberspace sovereignty*' – yet another confusing term for us to note, but essentially meant to encompass *anything* in the digital realm² (Schia & Gjesvik, 2017).

The second data sovereignty concept arises from the literature of indigenous peoples and views data as a resource, which must be treated in the same way as other products of cultural knowledge³. This definition is clearly quite different to a geopolitical one, by focusing on sovereignty as applied to communities or individuals. Whilst the geopolitical focus has enabled governments to control transborder data flow, it has done little to consciously advance individuals' access to, and control of, their own data. That, of course, was never the goal. Indigenous data sovereignty seeks to gain control over data as a *right*, from the position of a colonised people. Te Mana Raraunga⁴ have done more than any group within Aotearoa New Zealand to raise the profile of this non-geopolitical dimension of data sovereignty. I will argue that their goals and aspirations, particularly when supplemented with novel technological means to practically enable data sovereignty, offer significant benefit for everyone. Although, in saying this, I certainly do not wish to downplay the lived experience of systemic inequity that underpins that movement.

What this should indicate is that the term 'sovereignty' actually has layers of meaning, and this becomes even more evident when applied in the realm of data. Classically, of course, 'sovereignty' was understood firmly within the bounds of political theory and international law, for example in the work of Thomas Hobbes on 'social contract' theory in 1651 (Hobbes, 1962). The modern use of it by Māori also links to this understanding via constitutional law in the Aotearoa New Zealand context, the relevant vehicles being the 1840 Treaty of Waitangi and the Treaty of Waitangi Act 1975. Nevertheless, language evolves and meaning shifts.

Pasifika are also building knowledge and capability around their data sovereignty, and have asserted rights over data as a principle of *natural justice* in the absence of a legal document such as the Treaty of Waitangi (Moana Research, 2021). Similarly, the notion of sovereignty is found in contemporary radical feminism where, for example, Whisnant considers that it is the "systematic deprivation of bodily sovereignty [that] defines the oppressed condition of women"

¹A wide range of countries have data localisation legislation of differing types in place – for example, China, Spain, Canada, Indonesia, South Korea and Nigeria.

²As such, it is intended to be much broader than mere 'data sovereignty'.

³In te ao Māori it is considered a *taonga*, which can be translated as a "treasure" or something precious. This is because it is culturally significant and indicates a connection to Treaty rights (Sporle, Hudson, & West, 2021).

⁴The Māori Data Sovereignty Network (<https://www.temanararaunga.maori.nz/>).

2. Literature review

Table 2.1.: Types of data sovereignty. Adapted from <https://www.temanararaunga.maori.nz/> (accessed 20 April 2022).

Type	Description
Data sovereignty	Data is subject to the laws of the nation within which it is stored
Indigenous Data Sovereignty	Data is subject to the laws of the nation from which it is collected
Māori Data Sovereignty	Māori data is subject to Māori governance

(2010, p. 161). Finally, we have already noted the use of ‘cyberspace sovereignty’ in China’s 2017 Cybersecurity Law.

All of this is to state the point that ‘sovereignty’ is already being used by disparate groups to describe feelings of powerlessness or injustice. I should not, therefore, be afraid to further extend the notion that individuals and groups should have decision-making authority over their data, and to refer to it as a type of sovereignty.

But, for this thesis, I do find working definitions of data sovereignty – found within the indigenous literature – to be lacking. For example, the website of Te Mana Raraunga provides a typology which is shown in table 2.1. In it we can see that the first two definitions align more closely with a *geopolitical* understanding of data sovereignty; they rely on the prevailing legislation in either the jurisdiction of data residence, or jurisdiction of origin.

In either case, this may be problematic. To take an extreme example using this framework, it would be difficult to claim that a Russian citizen has sovereignty over their data stored within the federation. Similarly an ethnic Hmong living in China, who stores personal data in Australia, would arguably be *worse* off under the ‘Indigenous Data Sovereignty’ definition. Of course, we cannot simply ignore international law and regulation and I do not claim it to be unimportant. My concern is that – like the unreflective dominance of the centralised hegemony – thinking only within its boundaries can constrain our understanding of possibilities.

If the goal of Te Mana Raraunga is simply *governance* over Māori data⁵, then I propose a ‘radical data sovereignty’ which utilises the following definition:

The removal of centralised data silos, and the transition of data into the hands of individuals, groups and communities to be utilised as best meets their needs.

⁵Here I acknowledge my unwarranted privilege as a Pākehā in appearing to criticise the importance of governance. My intention is simply to highlight the fact that distributed and decentralised technologies have the potential to take us *much further* than this.

2. Literature review

Legal and regulatory frameworks are certainly important (and will be discussed in more depth later in this chapter), but they do not form part of my working definition. I am positioning the ability to move on from centralisation, that the nascent distributed technologies offer, as the focus in this research. In this way, ‘radical data sovereignty’ is enabled by ‘radical distribution’ and this last concept is discussed further in section 2.1.4.

2.1.2. Data privacy

This concept is included here only because of a tendency to conflate language around security, privacy and related concepts (Wahlstrom, Ulhaq, & Burmeister, 2020). They are certainly discrete concepts, but tend to interplay and overlap closely in the realm of data and digital services. For example, centralisation has the perceived security advantage that it is a well-tested and familiar paradigm. By the same token, however, it introduces a major potential privacy vulnerability by virtue of it storing large amounts of data in a single place. Distribution of data has the potential to avoid some of the most prevalent practical and governance risks around security and privacy.

If privacy truly is about “the right to be let alone” (S. Warren & Brandeis, 1890)⁶, then what better way to express that than being in complete and sole control of your personal information? I will briefly review some of the interplay between privacy and sovereignty to draw out some of these issues.

Before the computer era, there have long been provisions to ensure that posted letters should only be opened by the addressee⁷. In Aotearoa New Zealand the Privacy Act was refreshed in 2020 to better align with advances in technology. However, legal tools often sit uncomfortably alongside modern technologies which route traffic across the world in multiple steps and, in some cases, are openly used by advertising companies to mine personal information so it can be used further or sold. In the case of email, for example, Google have been criticised for scanning emails of users of its Gmail service (Yurieff, 2018), something that has only become more embedded with the advent of AI-assisted ‘smart composer’ features. Similarly, Facebook has been criticised for the monitoring and security of its own Messenger service (Vogelstein, 2018), which remains unencrypted by default.

Many legal writers have been despondent about the prospects of ever capturing a universal, or even broad, legal definition of privacy. Post, for example, has complained that privacy is “so complex, ... so engorged with various and distinct meanings, that I sometimes despair whether it can be usefully addressed at all” (2001, p. 2087). Despite this, others contend that dealing

⁶Although, certainly, there are many competing interpretations and these have proliferated in the digital era.

⁷This is generally derived from the Fourth Amendment of the US Constitution, for example, but has explicitly been part of Title 18 of the United States Code since 1948, making it a federal offence.

2. Literature review

with privacy in the context of specific claims actually has been entirely possible and, in fact, has worked satisfactorily (Penk, 2016b). This means, however, that national privacy law in a given jurisdiction can often be complex and fragmented and without a clear overarching legislative framework.

This can be seen in Aotearoa New Zealand where the government felt it necessary to pass legislation (the Intelligence and Security Act 2017) to unify the fragmentation that had occurred solely in respect of intelligence and national security services. Aotearoa New Zealand still has a complex patchwork of privacy legislation and regulation across other domains, and this is discussed in more depth in section 2.2.5.

Notwithstanding the above problem of *defining* privacy, it is true that the modern technology landscape now provides “unrealised opportunities to challenge [our] privacy expectations” (Harvey, 2016, p. 386). Even since 2016, many such opportunities have indeed been realised (Citron, 2021; Doberstein, Charbonneau, Morin, Despatie, et al., 2022; Herskind, Katsikouli, & Dragoni, 2020) thus making the case for relevant and robust data privacy legislation ever stronger. However, it should also be noted that there is a cultural component and expectations around privacy are not universal or homogeneous – this speaks to the importance of *social context* in understanding privacy (Ribak & Turow, 2003). Toh, for example, complains that any need to focus on privacy in the Singapore context is simply a result of “the unfortunate influence of American and European culture” (2016, para. 17). Similarly, in the indigenous peoples’ context, “notions of property, ownership and privacy (Western in origin) are foreign to a normative and social system that emphasizes totality and interconnectedness” (Williams, Vis-Dunbar, & Weber, 2011, p.22).

Garett and Young (2022) consider an additional nuance, which is the propensity for attitudes and perspectives to change over time and to be modified by the social context. For example, their US survey found that 75% of respondents who had been unwilling to share health data for research had completely changed their attitude one month after the COVID-19 pandemic begun. The urgency of a public health event, presumably, encouraged people to be more relaxed about the sharing of data under an assumption that benefit would accrue to people who needed it. This explanation was not tested by the authors, but it is a reasonable guess based on the literature we have already reviewed.

The “right to be let alone” requires the ability to completely remove oneself from any intrusion (Penk, 2016b). In the context of data in the modern era, this becomes more complex where it is necessary to provide electronic information to carry out routine activities, and where it is possible for people to identify others via social media. The technical ease by which data can be merged and joined has now led not only to greater private sector data matching and opportunities for surveillance (Biddle & Poulson, 2022; Kang & Frenkel, 2018), but also to similar efforts by governments. Whilst such matching can be justified on grounds of efficiency and the opportunity

2. Literature review

for better resource allocation, and indeed has proven effective in identifying risk, it is also true that it opens further complexity to the data privacy debate (Garvie, 2020; Hurley, 2018).

The problem of privacy in the modern era is neatly summed up by Lepore, thus:

“The defence of privacy follows, and never precedes, the emergence of new technologies for the exposure of secrets. In other words, the case for privacy always comes too late. The horse is out of the barn. The post office has opened your mail. Your photograph is on Facebook. Google already knows that, notwithstanding your demographic, you hate kale” (2013, para. 6).

True privacy means that an individual does not have to reveal any information, directly or indirectly, if they do not wish to. Centralisation has a huge influence on this because, by definition, your information is stored under someone else’s control. In Aotearoa New Zealand the penalties for privacy breaches were increased under the Privacy Act 2020 but, reflecting Lepore (2013), if you do ever find out that your data has been used in a way that you object to, or do not consent to, it will always be *after the fact*. Put simply, distributing data removes this problem entirely by making data holders (individuals, groups or communities) the decision-makers about what happens with their data⁸.

2.1.3. Centralisation

There is in fact a wealth of literature around the concepts of centralisation, decentralisation and distribution, if we look outside the narrow confines of this thesis. For example, organisational and management theory was engaged in existential debates about the relative value of each throughout the 1980s (Peters & Waterman, 1982). This thinking then goes on to influence discussion throughout the 1990s around the rise of ‘New Public Management’ and how democracy might be refashioned to give support to its aims (Burns, Hambleton, & Hoggett, 1994). In the IS context, however, this debate (if there is a debate) is very new.

Where organisational theory and public policy ostensibly have the ontological freedom to debate any conceivable position, IS has hard technical limits which constrain it. Clearly, IS and Computer Science do not operate solely in the realm of what is currently possible, or else we would not today have Artificial Neural Networks or the dubious benefits of facial recognition. A hard technical limit can be surpassed, as long as there is sufficient motivation and a compelling use case.

We have already noted that the development of relational databases in the 1970s was a paradigm shift in computing. It permitted extremely efficient storage and retrieval of information, as well

⁸There are of course many contextual nuances to this, which will be explored further in section 2.3.2.

2. Literature review

as a dependable and consistent knowledge model surrounding it. Even in the modern era, where scalable and distributed cloud computing services have increased flexibility, almost every web-site and software application today relies on the centralised model.

The general concept may be obvious, but we should be clear – under centralisation, all data flows from the outside (from users, or devices, or other web services) into a central point. This might be a relational database, or a document store or a cluster. There are many different approaches and so, whilst the technical layer certainly indicates and encourages centralisation, we should actually be thinking about it as more of an *implementation paradigm*. For example, Facebook has more than 1.22 exabytes of personal data. This will not be stored in one relational database, or in one location, or in one format. The point is that Facebook’s data – provided willingly by users of its platform – is under Facebook’s control and, subject to the legislative and regulatory framework in the jurisdictions where that data is located (geopolitical data sovereignty), they can more or less do with that data as they please.

Facebook users were not asked if they wanted their personal data to be shared with Cambridge Analytica; Twitter users are not asked if they want their location information to be provided in real time to Anomaly Six (Biddle & Poulson, 2022). So, centralisation as it relates to personal data is almost a philosophical position whereby a company has the technical means to receive data and then use it for some purpose. The internet trope is “if you do not pay for the product then you are the product”, and it is easy to see this at work here. Facebook and Twitter are large for-profit corporations; they require a business model to be sustainable, and the only product they have is a wealth of personal data which is of great interest to advertisers, marketers and cyber-arms/surveillance companies.

Outside of the for-profit space, governments and public sector agencies also rely on centralisation. In the Aotearoa New Zealand health sector we actually see a hierarchy of centralised data silos, as shown in figure 2.1. In this representation we see, on the far left, a patient attending the Emergency Department at their local hospital where it is recorded in their computer system. This data is then subsequently replicated across three additional and completely separate data stores, administered by different entities.

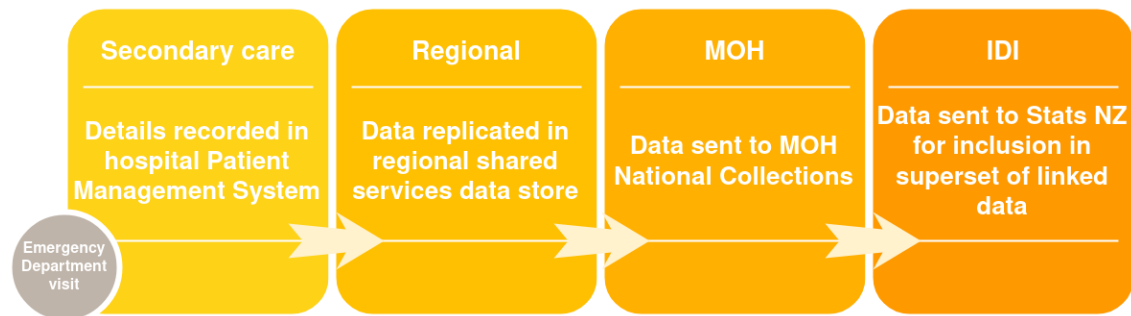
So, centralisation is not only about data being under the control of one entity; even in the public sector your data can be distributed across a number of different entities, each administering their own centralised data store. This is perfectly legitimate, of course, and, in the case of health data, may well be a pattern that the general public support⁹. I point this out here merely to reinforce the notion that centralisation is less a technical pattern than it is a business model or an implementation paradigm.

But what is common across centralised patterns in the public and for-profit sectors is a sense

⁹I will investigate this further in chapter 4.

2. Literature review

Figure 2.1.: Sample data flow across healthcare entities



that the individual, who is the subject of that data, can exert very little authority over what happens to that data and, in many cases, will not have any cognition of how it may be subsequently shared or used.

2.1.4. Decentralisation

Ambiguity around, and misuse of, the term ‘decentralisation’ is not a new phenomenon. In reviewing the implementation of New Public Management in Finland, Yliaska (2015) refers to a “battle” for the true meaning of decentralisation. So it may not surprise us that there is likely to be ongoing debate about what it means today in the context of personal data.

Decentralisation is not a new concept. Classically it has been utilised extensively in political science and political economy. For example, de Tocqueville (1898) considered administrative decentralisation as being essential for healthy democracy, by virtue of its ability to engage people in civic duty. We therefore see very early on that this concept has nuance and can be understood across different dimensions.

At present the term is being heavily utilised by proponents of blockchain and cryptocurrency. For example:

- The Ethereum website states on its main page that “Ethereum’s decentralized finance (DeFi) system never sleeps or discriminates”¹⁰
- Cardano claim that their product “combines pioneering technologies to provide unparalleled security and sustainability to decentralized applications, systems, and societies”¹¹
- Solana claim to be “Decentralized and unstoppable. Not only is Solana ultra-fast and low cost, it is censorship resistant. Meaning, the network will remain open for applications to

¹⁰As per <https://ethereum.org/en/>, accessed 25 April 2022.

¹¹<https://cardano.org/>, accessed 25 April 2022.

2. Literature review

run freely and transactions will never be stopped”¹².

The enigmatic and potentially fictional creator of Bitcoin, Satoshi Nakamoto, only stated in the original white paper that “there is no central authority to issue them [coins]” (Nakamoto, 2008, p.4). Does the absence of a central authority automatically result in decentralisation? It would certainly appear so, at first glance. Cryptocurrencies operate outside of the mainstream regulated financial markets (although, increasingly, they are being regulated). Thus, the *centre* is certainly not managing them. But, at the risk of entering a semantic rabbit hole, shedding one form of centralisation can potentially lead to merely differing forms of centralisation. For example:

“One of the key tenets of crypto is that it’s decentralized, as in no single person, company, or government is in control. But that may just be a ruse, as we’ve seen the recentralization of so much of the theoretically decentralized stuff” Dr Catherine Flick (Reader in Computing and Social Responsibility, De Montfort University), quoted in (Dailey, 2022, para.6).

A good example of this is the Non-Fungible Token (NFT) marketplace. As with cryptocurrency, NFTs utilise blockchains to universally validate ownership. This is done in a distributed fashion across all the nodes that are part of that network. However, the only way to actually buy an NFT is to go to a marketplace – centralised web stores which operate as listings aggregators so that NFTs can be searched or made available for sale. Currently, the dominant marketplace is OpenSea which, as at August 2021, was estimated to have more than a 90% market share. OpenSea is a private company that makes business decisions about what can happen on their platform. With a 90% market share, it is all but impossible to purchase an NFT without using OpenSea. How can we reconcile this with crypto’s promise of decentralisation?

Gottsegen (2021) makes the interesting point that — even where the computing or infrastructure layer is decentralised — *market* structures will tend to be centralised. This reflects a pattern we have seen in technology where certain players tend to become dominant – from Microsoft, in the realm of computing and software products, to Amazon, in the realm of e-commerce. The same pattern, indeed, occurs in cryptocurrency as well, where purchase or conversion of currency must be done via an exchange. In practice, these are centralised entities that may further enforce soft centralisation and control via withdrawal policy or cryptocurrency wallets.

But we should not forget the relevance of social context (Hoffman, Ibáñez, & Simper, 2020). For a cryptocurrency investor who is opposed to government regulation of money, then Bitcoin or Ethereum may well feel suitably decentralised. The zealot who insists that “no single person, company, or government is in control” will feel very disappointed about use of the term once they look more closely.

¹²<https://solana.com/>, accessed 25 April 2022.

2. Literature review

In this thesis I have the great liberty of defining, within reason, decentralisation as it relates to my research. Whilst my original writing did in fact focus on this concept as the means by which data sovereignty could be achieved, further research led me to Baran (1964) who, I will argue, conceptualises decentralisation very accurately as it currently relates to blockchain and cryptocurrency. That is, the removal of a single centralised arbiter which is replaced by multiple replicated points of centralisation. This model requires transiting centralised nodes in order to conduct transactions, even where a pre-existing centralised layer has been disaggregated. We will explore this further in section §2.3 where we will also identify inherent issues with this model, such as the risk of *re-centralisation*.

2.1.5. Distribution

Baran (1964) conceptualises a ‘distributed’ architecture as one in which all nodes (users) can connect directly with each other. There is no centralised authority, and no requirement to transit communications through any intermediary. In the context of personal health data this model is, to put it informally, as radical as it gets.

In the previous section we noted that recent cryptocurrency marketing has heavily leveraged the decentralisation concept. Popularly, this has been interpreted to mean that there is *no* central authority. In the case of blockchain this is demonstrably untrue, as we will see. Only in a distributed model can we confidently state that *all* centralisation has been removed, and users can connect directly with each other. Confusing things yet further, I should note that a core component of all blockchains is what is known as a ‘distributed ledger’. It would be more accurate to refer to this as a ‘decentralised ledger’. Conversely, technologies which are architecturally truly ‘distributed’ will also use the term ‘decentralised’ – perhaps because of its popular acceptance resulting from the cryptocurrency boom. For example, the Interplanetary File System (IPFS) refers to technical components as “decentralized”¹³, and Scuttlebutt describes itself as “a decentralised platform”¹⁴. I mention all this simply to highlight the semantic ambiguity in play. So then, what is ‘distribution’?

Baran’s framework (on which I rely throughout this chapter) visualises the distributed model as a mesh where all nodes can interoperate and connect. Where nodes are analogous to users, this is most easily understood as a peer-to-peer paradigm. That is, users connect directly with each other. Peer-to-peer architectures are not new. As far back as the late 90s, in fact, the model was popularised in the file-sharing scene by products such as Napster, Gnutella and Limewire that featured a centralised directory of available files, from which users can choose to download after connecting directly to the user who is offering the file. In this implementation a centralised

¹³<https://ipfs.io/> (accessed 7 May 2022).

¹⁴<https://scuttlebutt.nz/> (accessed 7 May 2022).

2. Literature review

component performs the function of indexing files and connecting users who want to share.

This became obsolete in the early 2000's, with the advent of BitTorrent. Still used widely today for a range of legitimate purposes, it represented a paradigm shift in peer-to-peer technology by implementing a Distributed Hash Table (DHT) for the first time. This enhancement enabled large numbers of users to simultaneously connect and share 'chunks' of data at very high speeds and with tremendous efficiency. This concept is also present in some nascent technologies that are categorised under a 'decentralised' umbrella, but architecturally are very strictly *distributed*. These are the technologies that are of primary interest in this thesis, since I argue that only the maximal removal of centralised components can lead us to data sovereignty.

2.2. Perspectives on data sovereignty

In this section, I will review a range of differing perspectives on data sovereignty. While I have already outlined how I will utilise the data sovereignty concept in this thesis, it is important to understand the knowledge base. Firstly we will identify two key ways of understanding data sovereignty – the geopolitical context and the indigenous context. We will then review the literature on what individuals think about it, before looking at some interesting international examples.

2.2.1. Geopolitical data sovereignty

A focus of much data sovereignty literature is the national security context – where data comes from and the jurisdiction in which it is stored. This has become a growing issue with the popularity of cloud computing and 'Software as a Service' (SaaS). There is an internet trope that "the cloud is just someone else's computer", which is technically correct. However Branscombe rightly notes that this is disingenuously reductive and offers a more complete, but less memorable, alternative:

"[The cloud is] a hyperscale, automated computer farm run by someone who's better at automation and security than you, and can buy electricity and servers and network connectivity more cheaply than you." (Branscombe, 2016, para.18).

Cloud data storage is always stored in a specific geographical location; the 'hyperscale computer farm' physically exists *somewhere*¹⁵. The laws of that land generally provide jurisdiction and

¹⁵For users based in Aotearoa New Zealand, the closest public cloud platform offerings are in Australia where both Microsoft and AWS have data centres located in Sydney and Melbourne. Both organisations have announced plans to build data centres within Aotearoa New Zealand by 2024.

2. Literature review

legal rights over that data, no matter where the actual end user is located (Irion, 2012; Kushwaha, Roguski, & Watson, 2020).

This growing awareness has led to several high-profile examples of attempts to reassert nation-state sovereignty, where there may otherwise have been ambiguity around who actually owns or controls the data. Amore (2018) notes that a US government programme announced in 2014¹⁶, whilst touted as an attempt to link government data sources, actually had border control, national security and immigration as its focus. It was a cloud computing programme about asserting the sovereignty of the US nation-state in its classical definition; that it is free from foreign interference¹⁷. Since that time, the US has initiated (and, in 2021, cancelled) a \$10bn Department of Defence cloud computing contract aiming to leverage cloud capabilities from data centres based in the continental US.

Another interesting example is provided by the enacting of Federal Law 242-FZ in Russia in 2015. Prior to 2015, Russia already had in place a “surveillance state ... that would have made the Soviet KGB (Committee for State Security) envious” (Soldatov & Borogan, 2013, para. 5). The national ‘Lawful Intercept’ programme – the System of Operative-Investigative Measures (SORM) – requires internet service providers¹⁸ to pay for and install equipment which permits real-time surveillance of telephone communications and internet activity (Roudik, 2016). SORM technology, and enabling statute, has also been introduced into Belarus, Kyrgyzstan, Kazakhstan and Ukraine (Soldatov & Borogan, 2013). Consolidating this, Federal Law 242-FZ requires that all internet services in Russia dealing with ‘personal data’ must be based within the territory of the Russian Federation. Notwithstanding the difficulty involved in easily distinguishing personal and non-personal data, the ambiguous definition of it in law meant that its provisions had far-reaching consequences. Its focus on personal data was executed without any protections or dependencies on individual consent, and thus effectively grants the state right of access to, and control of, personal data. Savelyev suggests that:

“it may look like excessive paternalism and intrusion in the private sphere of an individual ... [but] there is no feasible alternative to the mandatory nature of data localization requirements..., if governments want to ensure they work in practice and there are no economic or other incentives that could ensure compliance with them on a voluntary basis” (2016, p. 134).

Federal Law 242-FZ was introduced with no public debate or engagement, and enacted rapidly under the auspices of a national security agenda. The timing of the law coincided with critical

¹⁶Titled ICITE with a \$600m contract awarded to Amazon (Konkel, 2014), although very little information remains accessible.

¹⁷Or enabling the ‘right to be let alone’, if we wished to draw a link between privacy law and national security.

¹⁸Many foreign organisations, such as Cisco and Nokia, were also later found to have been instrumental in implementing SORM (Satariano, Mozur, & Krolik, 2022).

2. Literature review

geopolitical issues for Russia – not least the illegal annexation of Crimea, international sanctions and terrorist attacks carried out by Dagestani separatists. Similar to the US PATRIOT Act, 242-FZ was fast-tracked with no critical analysis or debate because “arguments appealing to national security are immune from almost all the possible criticisms in the time of crises” (Savelyev, 2016, p. 141), and the author further suggests that a key part of the agenda was simply to increase control of the internet. Bowman (2017) reinforces this point by suggesting that data localisation programmes are a double-edged sword – capable at once of asserting nation-state sovereignty over data, but also potentially allowing state actors easier access to the personal information of citizens. Additional uses of these mandates transpired during the illegal invasion of Ukraine in 2022, when the state communications regulator *Roskomnadzor* announced it had blocked access to Facebook and Twitter (Milmo, 2022).

‘Data localisation’ is also a feature of the technology policy landscape in China which, by virtue of its ‘Great Firewall’¹⁹, is already considered to be “one of the strictest countries in the world when it comes to the internet freedom [of citizens]” (Chandel et al., 2019, p.111).

The introduction of the 2017 Cybersecurity Law has divided opinion; something, again, that will be influenced by social context. For example, Yang and Xu (2018) view it as a meaningful attempt to unify a fragmented range of *Decisions*²⁰ by the National People’s Congress by providing greater protections around personal data and clarifying responsibilities. Human Rights Watch, conversely, labelled it a “a regressive measure that strengthens censorship, surveillance, and other controls over the Internet” (2016, para.1).

A range of articles in the law specifically aim to protect “national unity” and deter the incitement of separatism. Human Rights Watch (2016) notes that these terms – which are not clearly defined, but purposefully left ‘flexible’ – are used to punish peaceful protestors or pro-democracy activists²¹. This flexibility helps to “increase the government’s grounds to make wide assertions about the need for investigation” (Wagner, 2017, para.10) and restricts the possibility for those prosecuted under this law to contest, or defend themselves (Conger, 2016).

In terms of data sovereignty, though, the Cybersecurity Law also introduces the novel concept of ‘cyberspace sovereignty’. That is, it proclaims state sovereignty over the entirety of cyberspace that is within its capability to legislate or regulate. One of the most radical and

¹⁹A legislative and technological programme intended to provide the Chinese Communist Party with full control over the internet within China – a large part of which concerns blocking of foreign websites and services.

²⁰Specifically, the 2012 *Decision on Strengthening Protection of Online Information*, the 2012 *Notice Regarding Strengthening the Management of Network Access for Mobile Smart Terminals* and the 2013 *Decision on Amending the Law of the People’s Republic of China on the Protection of Consumer Rights and Interests*. These instruments have legislative weight within China.

²¹Indeed, ordinarily benign and ubiquitous software tools such as Virtual Private Networks – used commonly in a wide range of applications – have been branded “terrorist software”, simply by virtue of the ability to circumvent the Great Firewall (Lam, 2020).

2. Literature review

exciting facets of the early internet was precisely its ability to connect the world, and cross borders; China, in drafting this legislation, expresses a belief that cyberspace is synonymous with state sovereignty and national security. J.-A. Lee views this as the natural culmination of the regulatory trend in China over the last twenty years, very much entwined with its geopolitical positioning:

“As numerous internet regulations aim to restrict certain fundamental human rights, such as free speech and privacy, it is not surprising that China has extended its sovereignty to include cyberspace in order to obviate foreign interference.” (J.-A. Lee, 2018, p.69).

What does cyberspace sovereignty mean in practice? It can be seen as an aggregation of data localisation, the Great Firewall, oversight of all ICT infrastructure, and measures which reduce online anonymity. In short, China’s perspective is that the internet as a whole is something that individual countries should *control* (Lindsay, 2014).

Whilst we cannot ignore well-documented evidence of state-sponsored genocide within China (Gunter, 2021; Kanat, 2021), it will also be obvious that the language around sovereignty and freedom from foreign interference is quite similar across both the Chinese and US contexts. The other notable example we shall explore is that of the European Union (EU) which has consciously moved away from this rhetoric – ironically, perhaps, as a result of US activity.

The Edward Snowden PRISM leaks in 2013 prompted both consternation and a desire for member states to protect themselves from interference. The extent of US intelligence agency reach into data which foreign jurisdictions considered their own prompted concern primarily around the issue of *trust* across the EU (European Commission, 2013). Although interpreted by some as a move towards a separate European cloud (Amoore, 2018), the European Commission memo explicitly stated that “the Commission is strongly against a “Fortress Europe” approach to cloud computing” (2013, para. 4), and thus drew a clear distinction from the localisation initiatives already discussed, from Russia, China and the US. The EU focus was much more about competitiveness, incentivising adoption of cloud and restoring the public confidence which the PRISM leaks had harmed.

A critical enabler for reducing the barriers to the free flow of data is harmonisation of regulatory frameworks and privacy legislation across EU member states. This takes shape in Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) [2016] OJ L119/1 (GDPR) which took effect in mid-2018.

From a data sovereignty perspective, the GDPR provides for much greater control and regu-

2. Literature review

lation over management of personal data, along with significantly increased penalties for non-compliance, and it has been presented to EU citizens on this basis. The net result, however, will be that EU states, and the European Union bureaucracy, will actually have much greater access to personal data especially given the broad provision which permits derogation from prohibitions on processing categories of personal data where, Article 9 of the GDPR states, “it is in the public interest to do so”. The regulation does not define the ‘public interest’, or indicate any threshold above which these derogations can be utilised.

We can draw a perhaps unexpected parallel here back to China which, in 2021, implemented a Personal Information Protection Law (PIPL) – widely seen as an enhancement, or strengthening, of provisions found in the GDPR (Z. Tan & Zhang, 2021). Superficially, it appears to be the most far-reaching privacy legislation globally. However, other commentators point out that, whilst the law does technically include the state apparatus, the legislation is actually aimed at reining in the natural monopolies that have developed in China around big tech (Xiao, 2021). For example, the Chinese State Administration for Market Regulation imposed a record US\$2.8bn fine on Alibaba in 2021, citing monopolistic behaviour and concerns about social stability. This precedent has been seen as a response to the growing personal influence and wealth of big tech firm heads, such as Jack Ma (Thompson & Lockwood, 2021). Regardless, we can certainly see this development in a sovereignty/nation-state context. Even if the Chinese Communist Party (CCP) are more interested in keeping personal data solely within their control, rather than the private sector, it still reflects a sovereignty issue. It is just one that is, once again, heavily shaped by social context.

In terms of the GDPR, the parallel is that, while it does make significant improvements for how individuals can access and assert control over their data, its explicit genesis in the PRISM leaks (European Commission, 2013) indicates that it is also about asserting EU supra-national sovereignty over the data of its citizens. So this can in fact still be considered a kind of localisation exercise and, indeed, has been criticised for exactly that (Bowman, 2017).

The above examples demonstrate that data sovereignty can be used to assert geopolitical sovereignty over nation-state data ‘assets’ by strengthening government ownership and control of information in order to protect it from a perceived or actual foreign surveillance threat (or, in China’s case, a private sector threat). Such motives are based in long-standing principles of national sovereignty and statehood²². Indeed, Irion contends that data sovereignty is a critical public policy problem because “it is a crucial dimension of national sovereignty that presupposes the nation state” (2012, p. 42). Where control of information is lost, *governmental legitimacy* is also lost. This helps to explain the common importance of geopolitical data sovereignty across a range of politically heterogeneous jurisdictions.

²²As described in the ‘Social Contract’ theory of Hobbes (1962), for example.

2. Literature review

The sovereignty problem posed by cloud computing, notwithstanding its cost effectiveness and convenience, is that where government data is in an unknown location, or multiple locations simultaneously, it is difficult – if not impossible – to concretely determine jurisdiction. Indeed, Russian moves to enforce data localisation were primarily a reaction to fear that if the 2014 sanctions were hardened, and resulted in being disconnected from public cloud services, it would be difficult to assert ownership over it. The localisation rule thus ensures that Russian data will *always* reside in Russia and governmental legitimacy cannot be usurped. This approach seems natural from a legal perspective, where territoriality is fundamental to the principles of international law (Weber, 2010). In this context, no state can take action against, or seek redress from, another state without their consent, or as part of a mutual assistance treaty (Osula, 2015).

Although jurisdiction over data in a cloud service can be difficult to determine, Bradshaw, Millard, and Walden (2011) assert that doing so is crucial and data sovereignty is thus fundamentally about transnational jurisdiction. This appears to support the data localisation trends already discussed. Outside of this, neither can data sovereignty be asserted on a *property rights* basis, since data and information are intangible. Whilst there are established concepts in law for ‘data protection’ and ‘information privacy’, these do not provide a clear legal basis for data sovereignty either (Irion, 2012).

A further aggravating factor is the lack of alignment internationally, not only around legal and regulatory frameworks but in more prosaic respects such as definition of key terms. Scoon and Ko (2016) note that many Asia-Pacific countries have a formal legal definition for ‘sensitive data’, but Aotearoa New Zealand and Singapore do not. There are most likely many other examples, but the point is that where data is already being federated all over the world by its redundancy in separate data centres, there should be greater efforts to align regulatory tools.

In summary, geopolitical data sovereignty is utilised widely across a range of different political regimes. It is tightly linked with national security, and a desire for sovereign nations or supranational entities to assert their ‘right to be let alone’. This understanding of data sovereignty is about controlling where data resides, but there is another understanding which is maturing and gaining momentum.

2.2.2. Indigenous and Māori data sovereignty

The second main conceptualisation understands data sovereignty as a way of protecting cultural knowledge and assets and, depending on the jurisdiction, to exert rights enshrined in treaties or other constitutional instruments. The defining of personal data as a cultural asset which requires meaningful change in power and control structures, in order to exert decision making power, implies that the benefits of this movement will not only accrue to indigenous peoples. This is not to say they should be viewed equally, of course, because indigenous peoples have

2. Literature review

an evidenced and researched history highlighting the damaging consequences of colonisation. This impact is not felt directly by the ‘European’ New Zealander or Australian, and yet there is general demand for a shift in the power dynamic around data and it is indigenous researchers who are understandably motivated to raise the issue.

The research base for this particular strand of data sovereignty, outside of the indigenous context, is rather small. Weber (2010), for example, picks up the geopolitical thread but asserts that sovereign borders, and geographically based systems, have lost importance and relevance in the current technological landscape. In defining new concepts of sovereignty over data which reflect this, he advocates for a global approach which moves beyond the nation-state. Despite this, it becomes *more* important that groups who traditionally have not had sovereign power are included – NGOs and businesses are mentioned, but certainly we could consider a number of identifiable groups here in the same way (not least indigenous peoples themselves).

Following from this, Taylor and Broeders (2015) take an unusual middle path through the debate by highlighting how the actions of large corporations in developing countries are significantly impacting not just on government’s sovereign power, but also the representation and knowledge economy of citizens. A neoliberal governance trend is first identified, where the collection of data is shifted from central governments to distributed places “where power accrues to those who hold the most data” (Taylor & Broeders, 2015, p. 229). Where power is exerted and realised through the collection and interpretation of data, the superficially philanthropic involvement of large tech corporations in developing countries actually represents a fundamental shift in power.

Using examples of telecommunications giant Orange operating in post-disaster Haiti, or Facebook’s internet.org project across the economic south, Taylor and Broeders (2015) suggest that these ventures are much less about reducing inequities in access to technology and more about reducing inequities in access to the *specific services of that provider*. Furthermore, the vast amounts of data collected introduce a problem in interpreting and using that data properly from a remote position, and from a very different epistemic base.

Kukutai and Cormack (2021) note this trend around ‘surveillance capitalism’ – a term popularised by Zuboff (2019) to encompass the capture and abuse of personal data by advertising companies such as Facebook and Google, but also the opaque data trade with governments and state security apparatuses²³ – and agree that it represents the threat of embedding inequality even further.

The problem with the use and interpretation of data is extended by Bishop (2016) who identifies that there is a clear process for capturing and utilising data – an output, or piece of analysis,

²³For example, Kirchgaessner (2022) exposes the use of Israeli ‘Pegasus’ spyware by governments.

2. Literature review

can never arrive by chance²⁴. As Walter and Carroll note:

“Statistics are human artifacts and in colonizing nation states such numbers applied to Indigenous Peoples have a raced reality. Their reality emerges not from the mathematically supported analytical techniques they allow but via the social, racial and cultural standpoint of their creators” (Walter & Carroll, 2021, p.2).

For indigenous peoples, this means that they should be supported to engage actively in these processes and advocate for their own data needs – in the same way that other lobby and interest groups do. But, because effective representation over this process is so costly and time-consuming, it generally ends up being shaped by the people who collect and produce the data – governments or for-profit organisations.

As far back as 2004, the UN Permanent Forum on Indigenous Issues (UNPFII) recommended that indigenous peoples are actively involved in collecting and interpreting data. This reflects a large research base around bias in algorithms and automated decision-making, which was one of the factors that led to the so-called ‘right to legibility of automated decision-making’ introduced in the GDPR (Malgieri & Comandé, 2017). But the indigenous literature is going farther than this, and identifying that one’s epistemological position has meaningful impacts on how we understand and use data – and this is a problem where that is being done by government agencies, or corporations, who have no cultural understanding and produce inadequate conclusions (Walter & Carroll, 2021).

Without the kind of representation recommended by UNPFII, government policy defaults to reliance on existing administrative toolsets. These are always shaped strongly by the policy environment, which is itself shaped by the majority. For example, reintroduction of a data field for *iwi* in the Aotearoa New Zealand census in 1991 was primarily about the government’s desire to monitor effectiveness of a policy to devolve social services provision to *iwi*. Then, later, the 2018 census was found to be poorly implemented by government agencies, with quality of data on Māori so low that it was “unpublishable” (Sporle et al., 2021). Most government data, in the Aotearoa New Zealand context at least, is generally a “by-product of social and economic data collected for the total population” (Bishop, 2016, p. 297). Even where these data sets are considered to be good quality, the fact remains they were created and implemented for a specific purpose and hold fundamentally different objectives. The aim of indigenous data sovereignty is to ensure that data represents their status, knowledge and world view appropriately and also that, further downstream, their knowledge and data outputs remain under their control.

The result of good quality indigenous data – designed, built and understood in culturally appropriate ways – is that effective policy can be designed, resources can be properly allocated

²⁴This mirrors the Reflexive Thematic Analysis methodology we will explore in chapter 7.

2. Literature review

and we can hope for positive policy outcomes (Walter & Carroll, 2021). However, “on many occasions, the situation of indigenous peoples remains invisible within national statistics” (Davis, 2016, p. 25). Jansen (2016) describes how effective data exchanges, and Māori representation, has resulted in effective targeting of health services in Auckland to manage the prevalence of rheumatic fever. Extending this, other writers critique indigenous engagement with Western notions of quantification, which is “at serious odds with many indigenous ontologies and epistemologies” (Morphy, 2016, p. 100).

The importance of understanding and representing data can be shown in examples of differing world views between indigenous and non-indigenous peoples. For example, calculating welfare dependency ratios by age contains an assumption that age is related to economic engagement. Even prior to this, however, it assumes that capacity to earn money as a means to live, is a basis for a ‘normal’ economy. The notion that young people or old people are automatically dependent does not translate well to indigenous contexts where, for example, elders are the keepers of sacred cultural knowledge. Conversely, it hides the burden of unpaid work which disproportionately falls to women (Seedat & Rondon, 2021).

Even in non-indigenous populations, the formal economy and its associated systems of measurement is known to misrepresent core parts of society, also predominantly composed of women (Waring, 2018). Similarly, Western demographic traditions embed assumptions around units of measurement – a family is a heterosexual couple with children who live in a single household, for example – and these do not give voice to the traditions of indigenous peoples, as well as many other groups (Morphy, 2016). In short, data – within a centralised paradigm, at least – merely offers a different type of colonialism (Kukutai & Cormack, 2021).

Despite being able to sustain complex economic and social structures, colonisation forcibly replaced indigenous societies and knowledge models with Western colonial models – “any datasets that they themselves had not introduced and imposed were inferior” (Pool, 2016, p. 59). The legacy of this has an important bearing today, where indigenous cultural traditions and world views have been lost or marginalised, and are not represented in government data. Invoking the debate initiated by Taylor and Broeders (2015), Pool goes on to suggest that we should not allow the positive progress being made under indigenous data sovereignty initiatives to be stymied by their reliance on internet platforms which are “under transnational corporate rule beyond the control of indigenous peoples, or the polity in which they live” (2016, p. 72). There is thus a two-fold issue for indigenous peoples – that of asserting their own data sovereignty locally, and that of influencing policy to mitigate any transnational surveillance capitalism threat posed by multinational corporations or other actors.

In terms of the path forward for indigenous data sovereignty, we briefly looked at what I considered a limited or constrained ambition around Māori data sovereignty (as described by the *Tē*

2. Literature review

Mana Raraunga website) in section 2.1.1. The focus there is governance. Walter and Carroll (2021) extend this thinking and identify that governance is the means to operationalise indigenous data sovereignty, but *access and control of data* is an equally important counterpart. Two key things are needed to enable this: the ability to actually obtain data needed for governance (you can't govern data if it is not held by you), and; being able to own and control that data. Where a focus is simply on governance, Māori would be second-class citizens in a centralised system because decisions can always be overruled and the data is architecturally fixed within the colonial state apparatus. Kukutai and Cormack (2021) refer to this as a type of 'data dependency' and note that indigenous data sovereignty should be about new and self-determined data which can be owned, accessed and controlled by indigenous groups.

Recent indigenous data sovereignty literature is therefore taking on a more assertive tone about how to achieve their goals – Walter and Carroll talk about “inverting the central role of data dramatically” (2021, p.14), while Kukutai and Cormack (2021) call for a “radical alternative” to data dependency. So we have heard indigenous researchers talk about the need for ownership, access and control of data in order to enable effective data governance, which in turn operationalises data sovereignty. But how can this be achieved? Kukutai and Cormack conclude by suggesting that

“The as-yet unrealized potential is for a decentralized or distributed system that disperses data and power away from a central location or authority and puts Māori data in Māori hands” (Kukutai & Cormack, 2021, p.30).

This is exactly what I am proposing in this thesis. Distribution offers the practical/technical component which enables the dispersal of power away from centralised data silos. The literature indicates that achieving this would be an important component in realising Māori data sovereignty. My research focus is not Māori data sovereignty, but it is reassuring to confirm alignment of principles with a movement that has a real stake in making it happen. I will discuss, in chapter 5 and chapter 6, development of a prototype distributed system that will be the first opportunity to gauge this unrealised potential.

2.2.3. Individual perspectives

This section will focus on the literature around individual and public attitudes towards data sovereignty. I focus on data sovereignty because it is the desired outcome; distribution of data is simply the mechanism by which it can be achieved.

Positioning data sovereignty in a way that meets people's needs hinges on what they actually think. There are two main issues (and two accompanying sub-issues) in this area:

1. The term 'data sovereignty' is likely to have little resonance with the public in general

2. Literature review

2. The literature indicates that the general public currently suffers from two main problems in making informed decisions around personal data. These are:

- a) An overestimation of legal and regulatory protections available
- b) Difficulty in fully understanding the data landscape, and all the nuances of how personal data is gathered and utilised.

Although not framed explicitly in terms of ‘data sovereignty’, there is a wealth of literature internationally around public attitudes to more familiar concepts such as privacy and confidentiality. In order to make sense of the wide range of interlinking themes arising from a review of literature in this area, I have utilised the thematic framework presented by Aitken et al. (2016) but adapted to reflect the wide range of other factors that have been found. Whilst that work developed seven key themes, I have merged several and added two others that I consider to be distinct. Those themes are summarised as:

Table 2.2.: Public attitudes to sharing data – key themes, adapted from Aitken et al. (2016).

Theme	Description
Control and consent	The public value control of their data, and proactive seeking of consent
Trust and transparency	Public trust in an organisation informs public support for use of their data
Demographic/contextual variables	Views on data sharing vary widely within groups and across differing contexts
Anonymisation	Related to trust, individuals have a range of views on how anonymous their data can be made – and this heavily influences attitudes to data sharing

Whilst a theme of ‘anonymisation’ could easily be merged within the ‘trust and transparency’ theme, I felt it worth drawing out separately since much of the literature distinguishes levels of acceptance around data sharing by its perception as being anonymous or otherwise. The issue of anonymity, furthermore, brings into play some of the more subtle ways that data can be collected and used and, at the same time, speaks directly to more traditional government use of data.

2.2.3.1. Control and consent

The amount of control that individuals wish to exert over their data is inversely correlated with the level of trust they have in the organisation collecting or managing it (Aitken et al., 2016). Indeed, we will see in chapter 4 that GPs are very highly trusted and most people would be happy to share all their data with them in perpetuity. The public sector is considered to be more trustworthy than the private sector, and this is a direct product of where benefit is perceived

2. Literature review

to accrue. Individuals were found to be less trusting of private sector organisations in general, but specifically asked more questions about consent and wished to exert more control over data where research outputs were thought to be for private or commercial gain.

Aitken et al. (2016) found, in summary, that individuals desire control over their data, and this takes place across four dimensions:

- What data are collected
- Who has access to the data
- With whom the data are shared, and how
- What the data are used for.

The problem with these four dimensions is that – within a centralised paradigm – much of it is opaque, or not accessible to individuals.

Focusing more on the private sphere, Malkin et al. (2018) found Smart TV owners felt as if they lacked control of the data sharing inherent in such devices; many were even completely unaware that viewing history and other data were being collected. The issue of control is important here since, whilst Malkin et al. assertively conclude that “people do not want their data to be repurposed and shared” (2018, p. 9), those same people felt that they lacked control over what information is shared or how it is used. Considering that Smart TV data fits into the category that has been considered as most sensitive to individuals (L. Lee, Lee, Egelman, & Wagner, 2016; Naeini et al., 2017), this is a serious concern.

Consent is the critical mechanism for allowing the exertion of control, and is thus an essential prerequisite. The Data Futures Partnership (2017a) found that New Zealanders desired strong and detailed control over their personal data. This cannot occur without a consent model that is based on fully informing participants and which permits a detailed, ongoing, and iterative cycle of consent and re-consent – particularly when involving Māori and indigenous collectives (Reeves et al., 2022). Whilst Aitken et al. discuss differing preferences for ‘opt in’ or ‘opt out’ models across studies, they found “a clear preference for varied or flexible consent models which would enable individuals to set limits on their consent or to indicate particular preferences or objections” (2016, p. 15). This pattern is also reflected in literature reviews conducted by Garrison et al. (2016) and Moon (2017).

Other literature identifies a differing perspective taken by those actually conducting research, or managing data collection. Whilst Ormond et al. (2009) found that informed consent was considered to be difficult to obtain where the subject matter of the data was complex, Jao et al. (2015) found that researchers considered detailed consent models to be simply impractical. They concluded that only very broad consent is ever necessary – “broad consent can sufficiently

2. Literature review

support autonomy in informed consent by acting as a decision to allow *others to decide*” (Jao et al., 2015, p. 273). That is, broad consent provides control by fully *delegating that control to others* – which may well appear counter-intuitive to those advocating for control of their personal data. This view is affirmed by Sanderson et al. (2017) who found in a large multi-site US survey that whether the consent model was broad, tiered, opt-in or opt-out, mattered very little; the choice of differing models did not ultimately affect whether people actually gave consent or not.

However the consent is initially obtained, it is important that people can retain control over that to which they give consent. Even where individuals were perfectly happy to share their data, it remained important that they were proactively asked for consent. In India, by contrast, Hate et al. (2015) found that security and privacy are much more important to research participants than the nuances of seeking consent; implicit here is that keeping data secure, and preserving privacy, is acceptable even where full and informed consent has not been given. This is somewhat reflected in Jao et al. (2015), who found that researchers tended to consider consent as something to be balanced against potential benefits and the perceived risk of harm around any loss of privacy. For example, one researcher is quoted as saying “it is good for them to know, but personally I wouldn’t be very worried if what they didn’t know is not harming them” (mid-career male researcher, quoted in Jao et al. 2015, p. 269). This attitude is likely to be more prevalent given the findings of Reeves et al. (2022) that, despite a broad academic debate in the literature around sharing of data, only a small proportion took care to consider consulting with participants.

Regardless of these perspectives, there must be a responsibility for data agencies and researchers to gain *informed* consent²⁵, and this needs to be done in a range of ways to ensure understanding. Kass et al. (2015), for example, found that utilising a range of multimedia formats for different audiences was associated with a better understanding with which to underpin informed consent in clinical trials.

Reflecting the above, the Data Futures Partnership (2017a) found that New Zealanders want to see how their data is used, but also specifically to give permission for its use on a case-by-case basis. Concerns were raised that consent is often a one-off activity, and secondary uses may only come to light where there has been a breach or incident (reflecting the notion that the defence of privacy is always retrospective (Lepore, 2013), which I will contest only applies within a centralised paradigm). People wanted:

- Control over access to, and use of, data – “*I want to be asked, have the option of opting in and give permission each time*”, and “*each time I need to know who and for what, and give a digital tick*” (participants quoted in Data Futures Partnership, 2017a, p. 30).

²⁵In fact there is a legislative provision for this in many jurisdictions attached to differing data types. In Aotearoa New Zealand this is specified, for example, in Rule 3 of the Health Information Privacy Code 2020.

2. Literature review

- Proactive notification of changes – “*you should get a notification saying ‘here’s your current status - are you happy with this?’*” (participant quoted in Data Futures Partnership, 2017a, p. 30).
- Flexible permissions – “*I want to be able to give personal permission for access to specific things at specific points*”, and “*I want to have a choice about who can access the data and for how long*” (participants quoted in Data Futures Partnership, 2017a, p. 30).
- Access to an independent authority. Currently individuals have recourse only to the Privacy Commissioner, in the event of a breach of Privacy Act provisions, or otherwise to local privacy officers within an organisation.

The first three of these are all formally considered as requirements for design of the prototype artifact (please see chapter 5).

Of particular interest amongst these conclusions, is the finding that individuals do not have access to an independent authority which could effectively deal with a wide range of data use problems. Whilst the Privacy Commissioner does have statutory powers under the Privacy Act 2020, no body currently has any oversight over more subtle issues – such as improper interpretation of data, or algorithmic bias – where it may result in a negative or potentially harmful outcome. The Data Futures Partnership (2017a) found support for the concept of a ‘Data Use Commissioner’, which would provide an additional regulatory layer to that of the Privacy Commissioner. This would require legislative change but, more importantly, a clear conceptualisation of how personal data *should* be treated and what formal safeguards should be put in place.

Arguably, however, the need for such a role could be fulfilled by a carefully considered distributed data ecosystem. If groups were owning, controlling and interpreting their own data, then the sovereignty and *mana* that would be attached to that would probably surpass any benefit that could be offered by another statutory role – particularly, for Māori, where the statutory role is part of the colonial state apparatus rather than separate from it.

2.2.3.2. Trust and transparency

In draft data use guidelines, the Data Futures Partnership (2017b) firstly place ‘social licence’²⁶ as being a critical enabler in appropriate use of data, but also affirm the importance of transparency in *building* social licence. This is neatly summarised by Jao et al. who assert that public concerns can be significantly mitigated “with increasing understanding of the potential value and protections that could be implemented for data sharing” (2015, p. 268).

²⁶Defined as a generalised acceptance of data use, which can be built by following good practice guidance – with transparency at its core.

2. Literature review

Individuals' levels of trust in organisations (or governments, or private companies) dictate their support for use of personal data. Aitken et al. (2016), for example, found that public sector organisations are generally trusted more since they are usually subject to stringent accountability procedures and data protection regulations. There is a parallel mistrust of the private sector, which individuals perceive as being primarily motivated by profit and, hence, are less trustworthy. The distinguishing factor from the literature, however, was not simply whether an organisation was public sector or for-profit but where the benefits of the data usage will accrue.

Where use of data, and research, is perceived to be valuable to the public there is generally very favourable attitudes to data sharing – even where that occurs in the private sector (Aitken et al., 2016; Murad et al., 2017). Similarly, Belfrage, Helgesson, and Lynøe (2022) found that Swedish respondents were happy for data to be used *even without informed consent* in a health-care quality assurance context. That is, there is a perception that data use and analysis is used to directly improve quality of care and outcomes. As the data use case moves further away from this – clinical education and medical research, for example – informed consent becomes more important.

Belfrage et al. (2022) did, however, find very high levels of trust specifically in health care providers. This is often because individual GPs²⁷ are personally known and familiar to the subject. In terms of trust, therefore, it may be possible to build an abstract trust in a company, or sector, but individuals will always more easily place stronger trust in other individuals they interact regularly with in a data relationship – in this case, their doctor.

A 2015 survey of Guardian readers on this topic found that respondents were concerned about government's ability to safeguard their personal data. Despite this, respondents were even more concerned about *any* use of data by private companies, whilst at the same time admitting that they continue to provide data to those companies out of convenience (McMullan, 2015).

In Aotearoa New Zealand, research on trust found that participants wanted to understand the finer detail of organisational context around data sharing and use. Specifically, New Zealanders expressed a desire to know that their data will be:

- Used by qualified and competent persons
- Interpreted appropriately in context
- Secure, and subject to clear access protocols
- Protected by effective laws and regulation (Data Futures Partnership, 2017a).

²⁷Who, in Aotearoa New Zealand, are generally *for-profit* businesses with distinct business models from their Primary Health Organisation (PHO) who are often, but not always, non-profit.

2. Literature review

The issue around use by ‘competent persons’ certainly warrants some discussion – what does this mean? It is understandable that individuals want proper care taken of their data, but this becomes rather complex when one takes into account the full range of data usage – from collection, to storage, to management, to categorisation, interpretation, building into statistical models and, in case of government, its analysis and development into policy. At each of these points there are a myriad of potential actors, with their attendant myriad of attitudes, opinions, knowledge and experience. This is drawn out by Moon (2017) who notes that trust relationships are critical in the sharing of health data. A strong trust relationship is based on all of these touch points – from collecting only minimal data, to secure storage and acting in the patient’s best interests when sharing data. Nevertheless, participants realised that sharing data is an important part of an effective trust relationship in a healthcare context.

Trust in organisations is also very much contingent on the type of data concerned. Whereas individuals are generally willing to share non-sensitive data with trusted health providers (Garett & Young, 2022), Moon (2017) found that data on mental health, substance use, or sexually transmitted disease was strongly associated with lower levels of trust and willingness to share. Additional concerns found were that this information might be accessed by non-clinical staff, or it may find its way to private health insurers who would ‘penalise’ the individual (Reeves et al., 2022). Although affirming the importance of trust in general, Jao et al. assert that “trust in research seems to be independent of risks in practice” (2015, p. 274). This suggests that more can be done around transparency, and public engagement, to raise capability for assessing risk.

In a similar vein, other researchers have found in a healthcare context that, from the position of data ‘controllers’ (those who either collect or manage data on behalf of others), their biggest concern was fear of simply confusing patients or general anxiety resulting from negative media stories. GPs who were interviewed about how often they discuss use of data and data sharing with their patients revealed a wide range; 24% self-reported discussing it regularly, whereas 7% did not discuss it at all. Conversely, GPs held a perception that – a vast majority of the time – their patients simply accept data sharing, or just assumed it was already taking place. A ‘strong objection’ to data sharing occurred only in 9% of reported responses (Petrova, Barclay, Barclay, Barclay, et al., 2017). This aligns well with the findings presented by Garrett and Young, who reported that “fewer than 10% of participants were very uncomfortable with sharing data” (2022, p.2).

The Data Futures Partnership (2017a) noted that a key element in gaining trust was being reassured that data is secure and only accessed appropriately. This introduces two additional concepts to the trust issue. Firstly, trust is dependent on a perception that personal data will not be misused. There is a practical issue here, in that individuals must have some method of knowing that their data *has* been misused. Currently, privacy legislation in Aotearoa New

2. Literature review

Zealand only obliges organisations to notify the Office of the Privacy Commissioner in case of a privacy breach and, even then, only where there has been ‘significant harm’. Affected people need not be notified about the breach directly, if it is considered that informing them would cause more harm than not doing so. Secondly, it is important that individuals are fully aware of the legislative and regulatory framework within their jurisdiction, so as to properly assess risks and take privacy-enhancing measures where necessary. The issue here is that most people do not understand general principles around data collection and usage (Malkin et al., 2018) but, even when a risk is perceived, often choose to ignore it (McMullan, 2015). These two concepts will be discussed separately.

Potential misuse of data Aitken et al. (2016) found that concerns around potential misuse of data takes three key forms:

- Individual bad actors
- Commercial use or on-selling
- Political use.

Concerns around misuse of personal data is recognition of the fact that, particularly in a research context, data are managed, reviewed and analysed by *individuals*. It is entirely feasible that a research analyst, or database programmer, could misuse personal data in some way – there are certainly no technical or physical impediments to doing so. High profile media stories around hacking, data leaks, or accidental data breaches, have made the public cautious that security cannot always be guaranteed. In a Swedish study, Belfrage et al. (2022) found that the estimation of risk was also correlated with trust – further blurring the lines between these domains. For example, people with high trust in the health system tended to estimate a much lower security risk (either the likelihood or severity of unauthorised access to their data), than those who had low trust in general.

Study participants also expressed opposition to use of personal data for political purposes. Whilst this was framed in a context of objecting to use of data “for their [political parties’] own goals”, this is certainly a complex area. Perhaps the most simple example of this occurring is use of statutory administrative data in the US to redraw electoral districts. Even if “partisan gerrymandering has long been a facet of American politics” (Royden & Li, 2017, p. 3), the fact remains that reasonably basic administrative data is still used routinely to influence electoral outcomes (Lieb, 2017).

Outside of these three specific cases, Aitken et al. (2016) also note a general concern from participants around unanticipated *future uses* of their data, and also concerns around surveillance via data collection. Given that the data collection for studies which were included in their

2. Literature review

systematic review took place between 1999 and 2013, this is a well anticipated concern – public awareness of data surveillance via social media did not gain traction in the media until the Cambridge Analytica scandal in early 2018 (Kang & Frenkel, 2018).

Awareness of protections Other studies, focusing on less direct methods of gathering personal data, have noted that individuals tend to have an inflated perception of the legal protections available to them. This relates to the issue of trust and transparency, where internet users may assume that the very existence of a ‘Privacy Act’, or a website’s ‘privacy policy’, inherently protects them from data collection and reuse (Hoofnagle & King, 2008; Turow, Feldman, & Meltzer, 2005; Turow, Hennessy, & Draper, 2015). Using the reasonably mundane example of ‘Smart TVs’ we can identify that users are not informed enough about the technology changes present in a familiar and ubiquitous household item; its collection of data is far from obvious and there is no privacy policy to click (Malkin et al., 2018). Zeng, Mare, and Roesner (2017) found that even people with sophisticated ‘smart home’ implementations have very poor understanding of how these systems work and, thus, have a poor understanding of any potential security or privacy issues. This raises the issue of how proactive ‘data collectors’ must be to obtain informed consent but, more importantly for commercial use cases, what regulation is in place to make sure they are doing so?

Much of the literature around ‘digital literacy’ in this space draws out a more general fatalism around data use and the potential for control. Individuals continue to provide data, despite knowing it may be ‘misused’, either because it is convenient (McMullan, 2015) or simply because “you have no ability to restrict the amount of data that’s being vacuumed out of you” (Solon, 2018, para. 24).

This can be conceptualised as part of a ‘power and powerlessness’ narrative. NGO providers and users in Aotearoa New Zealand also exhibit a kind of “fatalistic” (Data Futures Partnership, 2017a, p. 36) mindset about data breaches, whilst a “narrative of powerlessness in the face of potential exploitation” (Hate et al., 2015, p. 246) underpinned attitudes to data sharing in India. In the context of consumer technology, Malkin et al. found that “despite strong opinions about data sharing, people will continue to welcome new technology ... our data indicates some level of resignation, as people reported that certain sharing relationships were likely to happen even if they were completely unacceptable” (2018, p. 8).

2.2.3.3. Demographic variance

One of the main conclusions that must be drawn, after reviewing the literature around public attitudes to data use, is that there is a great deal of complexity and heterogeneity. There are

2. Literature review

many bivariate or multivariate analyses of these attitudes in the literature, and these will help us understand how far it is possible to categorise them.

In a systematic review of 1,521 studies focusing on public attitudes to linking and sharing health data for research, Aitken et al. (2016) found that African Americans and LGBT participants voiced much stronger concern around misuse of data. Moon (2017) found that the demographic variables related to greatest disparity in data sharing preferences are: ethnicity, age, income and access to technology. Sanderson et al. (2017) meanwhile found that those individuals expressing greater concern about data sharing tended to be African American and those with lower educational attainment, with a slightly weaker correlation with self-reported religiosity.

In Aotearoa New Zealand, concerns in this area are centred around making sure that *communities* are fairly represented and there is a desire for data to be used more empathetically in local settings (Data Futures Partnership, 2017a). This has two key facets. One is the perception that data is ‘sent up’ to government, and then returns in the shape of aggregated analyses or pre-packaged policy advice. This is likely to have been done entirely in Wellington and may fail to have fully understood the complexity of the data itself, or the impacts of how such a policy might affect different communities (Kukutai & Cormack, 2021). The second associated facet is that ‘front-line’ providers feel excluded, or at best disassociated, from these processes, and do not have access to valuable information which would help them perform their functions – “it feels like the government sucks information from us, but never gives anything back” (participant quoted in Data Futures Partnership, 2017a, p.34). Furthermore, Aitken et al. (2016) note that individuals in their reviewed studies expressed a strong notion that they should not be conceived of simply as ‘sources of information’. The result is that communities feel unrepresented, and each community is composed of a unique diversity.

This accumulates to a general concern amongst individuals and community groups around potential use of data which lacks representation. Policies or decisions based on data skewed toward a majority group, rather than properly reflecting individual need, will leave anyone not part of that majority marginalised or disadvantaged (Hartz et al., 2011; Mezuk, Eaton, & Zandi, 2008; Walter & Carroll, 2021). This pattern also applies globally, rather than within a single jurisdiction. Research participants in low-to-middle income countries exhibit a keen awareness of global structural inequalities, and express a desire for any data sharing to be used to *narrow* those inequalities rather than widen them further (Hate et al., 2015; Jao et al., 2015; Sankoh & Ijsselmuiden, 2011). These sentiments begin to overlap with the theme of ‘trust and transparency’.

An additional factor which appears to consistently be important is age. The studies included in Aitken et al. (2016) showed a pattern of younger people being more concerned about privacy (and having a greater desire for control), while older people favoured less individual control

2. Literature review

and were less worried about possible impacts on privacy. However, it is important to note that participants from younger age groups were significantly under-represented across the included studies. Moon (2017) found a very nuanced range of attitudes amongst age groups. For example, whilst those under forty and those over sixty-five years of age tended to be *more open* to sharing health data, this was strongly mediated by other factors – such as understanding of the health system, or whether they had a chronic health condition. A qualitative study of adolescent (between 15 and 17 years old) attitudes to genetic research found that participants were supportive of contributing to scientific research, and “were unconcerned with a potential loss of privacy” (Murad et al., 2017, p. 935). This reflects the point made earlier about the perception of where benefit will accrue. However, this research also suggested that participants did not have enough knowledge about genetic research to make properly informed decisions. Overlapping with the importance of informed consent, Ormond et al. (2009) found that only half of those enrolled in a bio bank even understood that their DNA was being stored.

2.2.3.4. Anonymisation

“Data can be either useful or perfectly anonymous but never both” (Ohm, 2010, p. 1704).

The literature points to an inherent tension between prioritising confidentiality and deidentification on one hand, and ensuring that data can be linked and utilised effectively on the other. As already noted, it is very much associated with the theme of trust, since individuals tend to be more accepting of data sharing when they trust they will not be personally identifiable (Hate et al., 2015). Aitken et al. (2016) develop this by noting that anonymisation is very much a spectrum, and public attitudes to anonymity are difficult to interpret since it is hard to clarify exactly what levels and definitions of anonymisation are being used and, more importantly, what the public fully understand.

In Aotearoa New Zealand, the Data Futures Partnership (2017a) found a much clearer picture. They noted public scepticism that anonymisation was at all really possible in a small country like Aotearoa New Zealand – “my whānau is in [a small, remote town] with a small Māori population. If you state ethnicity and age I’ll know exactly who it is” (participant quoted in Data Futures Partnership, 2017a, p. 37). Whilst this intuitively seems more a problem for smaller countries or communities, other writers have demonstrated that this principle applies in any context. A respondent quoted in McMullan, for example, suggested that “anonymising data ... is useless because those with multiple data sets can easily cross reference them and undo the anonymisation” (2015, para. 27).

In a similar vein, other writers have highlighted key anonymisation issues from use of a range

2. Literature review

Table 2.3.: Anonymisation problems with OSN data

Data type	Information that can be accurately determined computationally	Reference
Content/text	Age and gender	Z. Wang et al. (2019)
Content/text	Mood of individuals	Eichstaedt and Weidman (2020)
Content/text	Lifestyle choices	Islam and Goldwasser (2021)
Content/text	Gender bias/sexism	Samory, Sen, Kohne, Flöck, and Wagner (2021)
Image	Location	Theiner, Müller-Budack, and Ewerth (2022)
Image	Relationship status of individuals	Choi, Budak, Romero, and Jurgens (2021)
Image	Age and gender	Abirami, Subashini, and Mahavaishnavi (2020)
Image	Sexual orientation	Kosinski and Wang (2018)
Location	Ethnicity/Cultural background	Wong, Zaïane, Davis, and Yasui (2020)
Location	Identification of users in a crowd	Gong, Daamen, Bozzon, and Hoogendoorn (2021)
Post metadata	Identity of user	Perez, Musolesi, and Stringhini (2018)
App metadata	User location / location prediction	Ochiai, Fukazawa, Yamada, Manabe, and Matsuo (2021)

of data types available via online social network (OSN) platforms, as summarised in table 2.3.

Outside of data obtained from OSNs, a large amount of data is collected and reused by governments. Although mostly using routine administrative data, we have already seen that this is being utilised to produce negative policy outcomes for marginalised groups (Royden & Li, 2017; Sporle et al., 2021).

How anonymous can an individual be in these data sets? Sweeney (2015) found that 87% of the US population can be uniquely identified from only three easily available census data points, whilst other research shows that 95% of people can be uniquely identified via only four location data points from cell phone towers (de Montjoye, Hidalgo, Verleysen, Blondel, et al., 2013). Finally Ohm (2010) has written extensively on the basic premise that anonymisation in data is simply *impossible*.

2. Literature review

If anonymisation is really so hard then it is also interesting to consider which data the public consider most sensitive and therefore, perhaps, most worthy of attempts at anonymisation. Both L. Lee et al. (2016) and Naeini et al. (2017), for example, found that photos or videos taken in a home environment are considered to be the most sensitive types of data by survey respondents. This is problematic considering the rapid adoption of ‘Smart TVs’, and the serious privacy issues associated with them already noted by Malkin et al. (2018). In differing contexts, both Belfrage et al. (2022) and Moon (2017) found that individuals are especially concerned about personal data where it deals with sexual health, mental health or substance use; particularly where there is a perceived risk it could be used to stigmatise them. This reinforces the indigenous data sovereignty literature, in a context outside of indigeneity, by reminding us that groups and communities are better placed to utilise data about themselves – rather than risk privacy breaches, security failures and analytical misrepresentation.

Given this, it is perhaps concerning to read about attitudes to data sharing within the research community which seem to infer that ‘anonymised’ data is inherently much safer and therefore requires fewer safeguards. Both Jao et al. (2015) and Ormond et al. (2009) found that researchers tended to consider aggregated or anonymised data as requiring less oversight than other types of data. Even ‘routine demographic data’ was considered to require less oversight, despite research showing how easy it is to identify individuals with only a few data points (de Montjoye et al., 2013; Sweeney, 2015).

Unfortunately there is not sufficient space here to conduct a full review of literature around applied digital re-identification methods, but the above table and discussion serves to provide a quick guide to what can be deduced from ubiquitous digital information. There are certainly many other examples of this anonymisation problem, but the intention here is simply to demonstrate the breadth of empirical evidence supporting public concern around anonymity.

2.2.3.5. Summary

The complexity of individual attitudes to data sharing makes it rather difficult to neatly tie up a conclusion. The Data Futures Partnership (2017a) made four key recommendations about how to progress the building of ‘social licence’ for data use in Aotearoa New Zealand; we have encompassed all of them in the foregoing discussion. In chapter 7 I argue that social licence is a product of *trust*, and Aitken et al. (2016) have noted that the level of demand for control over data is inversely correlated with levels of trust. But it is difficult to rely on extant research to accurately gauge levels of trust, since attitudes are changing all the time and none of the relevant research thus far has been framed in terms of data sovereignty, or the practical reality that data can be distributed. We are all constrained by the hegemony of centralisation.

2.2.4. International examples

In this section we will briefly review the data sovereignty and privacy landscape in four select jurisdictions. I do not intend to produce a canonical representation of global initiatives here; these four countries have been chosen because they represent different points on a multivariate continuum. For example, Estonia has produced a novel solution to geopolitical data sovereignty which differs markedly from that seen in authoritarian jurisdictions such as Russia or China. India has gone perhaps furthest in use of biometric identification, and embedding that within a society which is only recently catching up with global precedents around privacy. Uganda represent a model where superficially impressive achievements in legislation are found wanting and the state apparatus stands accused of human rights abuses. Finally, Taiwan has relatively weak privacy legislation but has built innovative tools around civic engagement, collaboration and transparency.

Put together, all four examples will elucidate the principles of data sovereignty in different ways. While no jurisdiction has gone down a path of distributed data as yet, these examples point to global trends and help to contextualise our understanding.

2.2.4.1. Estonia: E-government

The small Baltic nation of Estonia has been the subject of much media interest for its progress towards becoming a ‘digital society’, with Forbes gushing in 2021 that, “being the world’s one and only Digital Republic and “E-nation”, Estonia has given an entirely different outlook on what a truly human-centric society entails for digital citizens” (Minevich, 2021, para.6). The New Yorker has described it as “the most ambitious project in technological statecraft today” (Heller, 2017, para. 3) and the UK’s Government Digital Service noted in 2013 that Estonia is “probably the most joined up digital government in the world ... [where] ... the citizen is in control of their data” (Herlihy, 2013, para. 2).


‘E-Estonia’ is an umbrella programme, which aims to digitise as much of government as possible – both amongst government agencies and in all public-facing services. This initiative has its roots in the gaining of independence in 1989 and the ‘green field’ opportunity that Soviet withdrawal had created. Estonia resolved to make rapid progress in digital transformation such that, by the year 2000, it already had an impressive list of digital achievements under its belt²⁸ (Tikk, 2021). The centralised capture and management of data had become so sophisticated that, prior to the 2021 Census, there was serious debate about whether it could all be done via existing government data rather than polling its citizenry manually (Beltadze, 2020).

This last point should make it obvious that the Estonian model is based very much on a

²⁸For example a national e-health information system, a secure national data exchange layer and online elections.

2. Literature review

Figure 2.2.: X-Tee, the technical layer of E-Estonia (Republic of Estonia Information System Authority, 2022).



centralised paradigm; even the most impressive achievements today still leverage this venerable model, and they are no less impressive for it. But we should be very clear that, even where Estonian citizens can exert ‘control’ over their data (Herlihy, 2013), it is over a centralised data set that is managed by a government agency. The Estonian model does not help us to resolve the issue of ownership or sovereignty that has been so prominent in the indigenous literature. In fact, the Estonian Information System Authority (*Riigi Infosüsteemi Amet*), who are the state agency responsible for the digital government platform, are very clear in saying that “a state without data on its citizens is impossible in principle” (Republic of Estonia Information System Authority, 2019, para.13). Governments certainly do need data, but I contend that this does not mean it has to be *centralised*²⁹.

The heart of the ‘E-Estonia’ platform is an abstraction layer called “X-Tee” which permits data to be dynamically linked from disparate systems (a visual representation is shown in figure 2.2). Existing centralised databases are left in place. From the very beginning there was antipathy towards any notion of a centralised *superset* of government data such as Aotearoa New Zealand’s Integrated Data Infrastructure (IDI): “one of the biggest threats to the State is an uncontrolled centralisation of data into a single database” (Republic of Estonia Information System Authority,

²⁹I will discuss competing viewpoints on this in chapter 7.

2. Literature review

2019, para.15). This thinking was heavily influenced by the events of 1996, where a contractor to the Estonian police managed to link several state databases together and proceeded to sell the results on the ‘black market’ for very large sums of money, representing a privacy breach of enormous proportions³⁰. Centralised supersets of government data were therefore ruled out and X-Tee is designed to connect directly to relevant databases only when needed.

A supporting design principle was to avoid dealing with duplicate data and having to reconcile partial or overlapping datasets. If information is required about a person’s date of birth, then the requesting agency must connect directly with the population register and read it from there. Similarly, a person’s cancer diagnosis can only live in the cancer registry belonging to the Ministry of Health. These data cannot be copied and replicated in any other database. This has been made enforceable via legislative tools such as the 2001 Public Information Act³¹ and the 2014 Economic Activities Code Act³². By 2014, this ‘once only’ principle had become a formal initiative of 25 European Union member states and estimated to have a net impact of €5bn per year across the EU in efficiency savings (European Commission, Directorate-General of Communications Networks, Content & Technology, 2014).

The final key component was a national digital identification card. The card was first introduced in 2002, and the cryptographic signature it offers was made legally equivalent to a ‘wet signature’ some two years prior³³. The card has a wide number of uses, and is the means by which citizens interact with government digital services. The digital ID card is considered to be one of the key enabling factors in making X-Tee successful (Saputro et al., 2020). However, a flaw in the cards was identified in 2017 and nearly half of all ID cards were found to be compromised. The Estonian government’s response was criticised for not following best practice, and a legal case ensued between the government and the supplier (Parsovs, 2020).

Nevertheless E-Estonia goes much further than other examples of centralised data ecosystems, particularly in terms of the control and audit capability available to individuals (Vassil, 2015). Whilst this certainly does not provide *total* control, it is possible for Estonian citizens to delegate access to their Electronic Health Record (EHR) to other health professionals. By virtue of the EHR presenting essentially as a centralised data store, it is easy for health professionals to quickly find clinical information they need. Again, this is all made possible by the key enablers of digital identification, the data exchange layer and the application layer – this last one being

³⁰No English language articles exist, but a translated version of events from an Ekspressmedia article dated 17 April 2000 is available at https://epl-delfi-ee.translate.google/artikkel/50825984/puhka-rahus-andmebaaside-virtuoos-imre-perli?_x_tr_sl=auto&_x_tr_tl=en&_x_tr_hl=en&_x_tr_pto=wapp (accessed 26 April 2022).

³¹<https://www.riigiteataja.ee/en/eli/514112013001/consolide>.

³²<https://www.riigiteataja.ee/en/eli/530102013062/consolide>.

³³See Estonia’s 2000 Digital Signatures Act at <https://www.riigiteataja.ee/en/eli/ee/Riigikogu/act/524102016001/consolide>, although this was repealed in 2016 to align with superceding EU legislation.

2. Literature review

Figure 2.3.: Three elements of the Data Embassy Initiative. Sourced from Estonia Ministry of Economic Affairs and Communications (2016).



particularly pertinent in the health setting (Vassil, 2015).

Having identified Estonian initiatives to empower individuals with greater control of personal data, Estonia also presents some interesting developments around *geopolitical* sovereignty. The Estonian government commissioned a Microsoft research project in 2014, to understand how the extant initiatives discussed above could be utilised in the context of the public cloud. A primary concern was about maintaining public confidence, whilst at the same time using cloud services to protect and reinforce the digital continuity of E-Estonia. However, there was a recognition that using public cloud services introduced ambiguity around sovereignty, such as we have already noted in section 2.2.1. The project concluded by recommending a three part solution, which would see development of hybrid cloud models across existing Estonian government services:

1. Government Operated Cloud. The storage and maintenance of sensitive data within Estonian borders.
2. Physical Data Embassy. Data storage and maintenance of sensitive data at foreign Estonian embassy locations, or countries with appropriate bilateral agreements.
3. Virtual Data Embassy. Storage and backup of non-sensitive data in public cloud services (Estonia Ministry of Economic Affairs and Communications, 2016).

These points can be represented visually, as shown in figure 2.3.

2. Literature review

The Estonian government were ultimately not satisfied with the loss of sovereignty found in public cloud offerings. This was aggravated by large-scale Russian cyber-attacks in 2007, along with concern at Russia's illegal annexation of Crimea in 2014, which together raised the profile of 'true' geopolitical sovereignty within government. The government CIO at the time was quoted as saying:

"We didn't want to use a random cloud that's somewhere around the world, where we don't know what sort of rules and laws apply to the data we put there ... We wanted to have full jurisdiction over the data. No private cloud partner can really do that" – Siim Sikkut, quoted in Talmazan (2019, para.16).

Whereas countries such as Russia and China have approached this issue with localisation rules, Estonia took a very different approach by geographically distributing data onto Estonian diplomatic missions. A bilateral agreement with Luxembourg to host a digital embassy there (E-Estonia, 2017; Microsoft, 2017) was signed in 2017 and, while not physically in an Estonian embassy, the agreement does specify that the data stored in Luxembourg has all the protections and immunity of a diplomatic mission³⁴. The specific bilateral agreement is important in this case since Articles 22 (governing diplomatic properties) and 25 (governing diplomatic communications) of the Vienna Convention would not provide sufficient jurisdiction over data centres (McCluskey, 2015). Estonia therefore present a novel solution to geopolitical data sovereignty concerns, without resorting to localisation. Whilst Estonia has furthered this strategy with an additional digital embassy in South Korea (Selke, 2018), Rice sums up the significance of Estonia's digital embassy strategy thus:

"This is a type of cyber-sovereignty entirely different from the kind authoritarian regimes want to impose over their citizens. Instead of declaring a monopoly on information access, Estonia has used technology to expand its governance efficiency over its people (increasing the legitimacy of its institutions among its population) and to make those cyber-institutions durable in the case of territorial occupation, thus increasing the staying power of its sovereign claims" (Rice, 2019, para.15).

2.2.4.2. India: Biometric national identity and digital empowerment

Described as the world's largest biometric project (Dattani, 2020), India has separately implemented two large-scale programmes aiming to collect biometric information from its population

³⁴The full agreement can be accessed at https://www.riigiteataja.ee/aktilisa/2280/3201/8002/Lux_Info_Agreement.pdf, and notes under Article 5 that the data centre "shall be regarded as assets of the Republic of Estonia and shall enjoy immunity from every form of legal process".

2. Literature review

and to embed its usage into routine transactions requiring proof of identity. These two programmes are the National Population Register (NPR) and the Unique ID (now popularly referred to as *Aadhaar*, meaning ‘foundation’ in Hindi).

The NPR was created by a 2004 amendment to the Citizenship Act 1955, which mandated citizen registration and a National Identity Card. Its stated aims were to inform targeting of government services, to improve planning and to prevent identity fraud (Office of the Registrar General & Census Commissioner, India, 2012). Critics have stated however that the NPR’s origins are firmly rooted in national security and, specifically, a 1992 government policy (known as ‘Operation Pushback’) which aimed to identify and deport undocumented Bangladeshi immigrants (Sethi, 2011). The NPR was implemented at the time of the 2011 Census³⁵, and it was intended that the data gathered would form the basis for a national ID card, and accompanying national identification number.

The Aadhaar (Targeted Delivery of Financial and Other Subsidies, Benefits and Services) Act 2016³⁶ provided the legal basis to establish the Unique Identification Authority of India (UIDAI), and to issue the 12-digit identification numbers to each citizen present in the NPR. A critical differentiator from other national identity schemes, the UIDAI also capture biometric information about residents and this includes fingerprints, iris scans and photographs (as permitted under §3 of The Aadhaar Act 2016). Choudhary (2018) notes that India’s combining of biometric data *and* a national identity programme is unique amongst western countries³⁷.

With the explicit focus being on reducing friction in delivery of benefits and services (as per the title of the legislation itself), UIDAI have worked to ensure that the database can be accessed by agencies and suppliers. This is referred to as the ‘India Stack’³⁸, which comprises a framework of APIs and services which can be used to access and interoperate with Aadhaar data. At the time of writing, more than 1.3 billion Aadhaar numbers have been generated, and the system has processed more than 70 billion authentications in total (Unique Identification Authority of India, 2022). However, the integration of Aadhaar data with important service points, for example where pensions are administered or food rations are dispensed, has not been without criticism. There have been reports of the elderly being refused their pension (Khera, 2017), and children starving to death (Biswas, 2018), because documents holders are not part of Aadhaar. Meanwhile Dreze, Khalid, Khera, Somanchi, et al. (2017) found that the Aadhaar

³⁵Also notable by virtue of the fact it was only the second post-independence census to require information on caste (the previous occasion occurring in 1968).

³⁶Full text in English available at https://uidai.gov.in/images/targeted_delivery_of_financial_and_other_subsidies_benefits_and_services_13072016.pdf.

³⁷Moves towards this outcome in the UK were initiated by the Identity Cards Act 2006, which enabled a national biometric identity database. However this was repealed only four years later by the Identity Documents Act 2010, following public outrage and criticism from sector groups.

³⁸Please refer to <https://indiastack.org/index.html> for further information.

2. Literature review

biometric authentication points (fingerprint scanners) often became inoperable, due to technical or connectivity issues, and have resulted in people missing food rations or having their access to benefits terminated. Dattani (2020) contends that these consequences disproportionately impact already marginalised groups, rejecting the notion advanced by the government that technology can only be rational and neutral when applied to real world communities.

The Indian government's justification for Aadhaar was around benefit fraud, and "wastage". The original architect of the system talked about a vision where the state and the individual became closer, and where interactions were less traumatising. In order to allow marginalised populations equal access to the benefits and services for which they are eligible, more sophisticated identification systems were necessary – "those who most needed robust proof of identity in fact had the least access to it" (Nair, 2021, p.27). The narrative therefore, despite evidence to the contrary, is around equity and access. Indeed, Dattani is sceptical about this paternalism and notes that, rather than being less traumatising, Aadhaar has resulted in "the traumatic hardships of many who struggle to access welfare, or are excluded altogether" (2020, p.415).

Despite the fact that the legislation itself indicates (but does not specify) it is voluntary, its spread throughout Indian society has also been criticised. Deepalakshmi (2017) details more than 50 official schemes that require Aadhaar and they range from receiving benefits or pension, to opening bank accounts, applying for scholarships or receiving treatment for HIV. The ubiquity of Aadhaar has therefore meant that – whilst it was not mandatory to participate – it is a prerequisite to undertake many rather ordinary tasks, and is therefore as good as mandatory (Dass, 2011).

The release and development of the 'India Stack' has also been criticised. Whilst marketed with language around "financial inclusion" and "transparency", Dattani (2020) notes that it presents an unrivalled opportunity for the private sector to both identify and target large customer segments, and at the same time to gather large amounts of data around transactions, credit and spending patterns (Desai & Jasuja, 2016). In short, it has enabled a nationalised and state-sponsored form of 'surveillance capitalism' (Zuboff, 2019).

The petitioners in the 2017 Supreme Court case *Justice KS Puttaswamy (Retd.) & Anr. v. Union of India & Ors* argued that the practical reach of Aadhaar, and the scope of data linked to it, could be misused by Government and was an infringement of privacy. The Indian Government had previously been accused of not respecting the right to privacy, with its lawyers in this hearing arguing that Aadhaar should in fact be compulsory (Pandey, 2017). The Supreme Court ruling however stated that data privacy is an intrinsic part of Article 21 of the Constitution, which prevents the encroachment of life or personal liberty by the State (Biswas, 2017). This was a landmark ruling in India which, at that time, had no dedicated data protection or data privacy legislation.

2. Literature review

The Personal Data Protection (PDP) Bill, enacted in 2019, was seen as a response to the Supreme Court ruling and as an attempt to ‘catch up’ on global initiatives around the regulation of personal data. The Bill had very positive aspirations, stating that it intended to “create a relationship of trust between persons and entities processing the personal data”³⁹. Indeed, it went as far as introducing provision for the “right to be forgotten”, regulating data portability, and setting up a Data Protection Authority. Additionally, it was enacted as a key part of a broader ‘Digital Empowerment and Protection Architecture’ (DEPA), although this seems to be faltering⁴⁰. The PDP Bill received significant criticism, particularly after the government inserted exemptions for state agencies – with one senior judge declaring it posed the ability to turn India into an “Orwellian State” (Mandavia, 2019, para.1). However we should remember that exemptions on the grounds of national security are firmly in place within the GDPR and most other global privacy instruments; India is certainly not unique in this regard.

In summary, India’s narrative is around digital empowerment and protection of data. However, it feels as if all of these initiatives only occurred as a face-saving exercise following the 2017 Supreme Court ruling; little has been done to change the negative social outcomes already outlined. Dattani (2020) and Nair (2021), particularly, have voiced scepticism about these initiatives and contend that marginalised groups are now worse off in India than they were before, and that the technical solutions enabled by Aadhaar are being used to create a state-sponsored form of surveillance capitalism.

2.2.4.3. Uganda

The enactment of the 2019 Data Protection and Privacy Act (DPPA) made Uganda the first East African nation to pass any meaningful legislation specifically around data and privacy⁴¹. Clearly inspired by concepts found in the GDPR, the DPPA is based on constituencies of ‘data subjects’ and ‘data controllers’, who each have rights and responsibilities under the Act.

Firstly, it is interesting to note that a right to privacy is in fact codified in the constitution of Uganda, where it is formulated thus – “No person shall be subjected to interference with the privacy of that person’s home, correspondence, communication or other property” (Article 27, Constitution of the Republic of Uganda, 1995). Similar to many other jurisdictions, the DPPA contains exemptions for national security and intelligence operations. It does, however,

³⁹Preamble to Bill which can be accessed at http://164.100.47.4/BillsTexts/LSBillTexts/Asintroduced/373_2019_LS_Eng.pdf.

⁴⁰With accompanying language around commitment to “open source”, DEPA source code has supposedly been published at <https://github.com/iSPIRT/DEPA>. However there is no code at all in this repository and it has not been updated since October 2021 (as at 3 May 2022).

⁴¹22 other African countries in total have such legislation (United Nations Conference on Trade and Development – <https://unctad.org/page/data-protection-and-privacy-legislation-worldwide>)

2. Literature review

go somewhat further than this and also provides exemptions for “medical purposes” or “the performance of a public duty by a public body” (§7) both of which could be interpreted rather widely.

In terms of rights conferred by the Act it follows, at least in principle, much of what would be considered ‘best practice’ in the modern era, for example:

- Clarifying rights of children
- A right to explanation for automated decision-making
- Prohibiting any collection of special data such as religious belief or sexuality.

An interesting addition is the implementation of a ‘data protection register’, which will log all entities collecting or processing personal data and the reason for doing so. This must be publicly accessible under the Act, although the author could not find any evidence of this online at the time of writing⁴².

In reviewing the impact of this legislation, the Collaboration on International ICT Policy for East and Southern Africa (CIPESA) et al. (2019) were positive around the aims but considered that it did not have sufficient focus on strongly enforcing and regulating the rights of individuals. Indeed, “it is not to the standard of the leading European GDPR which offers significant guidance on personal data protection” (Collaboration on International ICT Policy for East and Southern Africa (CIPESA) et al., 2019, p.8). But GDPR is viewed almost as a reference point now for *best practice* privacy and data protection legislation; the problems in Uganda appear to go deeper.

Writing one year after enactment of the DPPA, Privacy International (2020) expressed concern that no meaningful action had been taken to ensure compliance with the Act – whether by state or non-state actors. A compliance scorecard prepared by the Civil Society Organisation (CSO) Unwanted Witness in 2021 similarly found a net compliance score of 35% against the DPPA (Unwanted Witness, 2021). Privacy International (2020) also express concern about some other components of the privacy and data protection landscape in Uganda. For example:

- An “intensification” of government data collection via the national identity card system.
- A Ugandan Police Force initiative to link a forensic CCTV system with national identity data (Joseph, 2021). This is made more alarming in the context of surveillance systems (purchased from China) which have been used by the Police to commit human rights abuses, including targeting anti-government protesters (Kafeero, 2020; Parkinson, Bariyo, & Chin, 2019).

⁴²May 2022.

2. Literature review

- Commissioning of a centralised government data centre to enact a form of state sector localisation (Monitor Uganda, 2020), which was criticised as a way of covertly connecting and sharing data outside of the reach of the DPPA (Privacy International, 2020).

What we see in Uganda, in summary, is a government that appears to be saying the right things about data protection and privacy. Other legislative instruments such as the 2011 Electronic Signatures Act, and the 2011 Computer Misuse Act, propelled Uganda to receive a 100% score for “protection of personal data” in the 2021 National Cyber Security Index⁴³. The DPPA has great promise, and clearly models itself on the GDPR. Yet, the criticism has been that government has not done anything to give life to it, and levels of compliance with its provisions are still very low.

We therefore have a hybrid environment where there is an impressive array of legislation and regulation, but very few benefits have accrued to ordinary Ugandans. In fact, some commentators argue that the government has an agenda which is almost antithetical to the aims of DPPA, and is based on tracking and surveillance and, in some cases, the obstruction of human rights such as those articulated in the Ugandan Constitution. A possible route forward for Uganda is to test this legislation in earnest. The judiciary has a track record of pushing back against homophobia in Uganda – for example, by striking down ‘anti-gay’ legislation in 2014, and invoking the constitutional right to privacy for people in same-sex relationships even when homosexuality itself remains illegal. As Mujuzi has noted:

“Although the executive and the legislature in Uganda take a hostile approach towards people in same-sex relationships, the judiciary has adopted an objective test in interpreting the Constitution and this has been a vehicle through which the rights and freedoms, such as the right to privacy, of people in same-sex relationships ... have been protected” (Mujuzi, 2012, p.118).

So the DPPA should become an important vehicle for testing government commitment to the protection of personal data, although by mid-2022 this has not yet occurred.

2.2.4.4. Taiwan

The independent and sovereign nation of Taiwan has attracted a lot of media attention for implementation of its “digital democracy” programme. This initiative may be viewed in a similar context to that of Estonia – a small nation with some sense of an existential threat from an aggressive and autocratic neighbour.

⁴³Please see <https://ncsi.ega.ee/country/ug/570> (accessed 28 April 2022). Uganda received a global ranking of 55, eight places ahead of Aotearoa New Zealand. Note that this same index gives India a score of 0% for protection of personal data.

2. Literature review

Taiwan's approach has been based on collaboration, citizen participation and transparency; an important first step was, in fact, not taken by government but by "civic hackers". In 2012 this saw the activist group build their own 'auditing system' for the government's budget. Data from government was taken and repurposed in a way that was interactive and easy to understand for ordinary citizens⁴⁴. Furthermore, it gave people a voice by permitting them to rate and comment on items (Tang, 2019). This became the g0v ("gov-zero") project, which has now expanded to offer citizen-built civic services such as:

- Hackfoldr – a place to organise documents and tools, to provide collaborative spaces for government agencies and communities
- Moedict – an improved version of the government online dictionary
- Political campaign finance transparency – A crowd sourcing tool to transform campaign finance data from its 'paper only' format, thereby making it transparent for all
- Legislator voting guide – A tool which tracks the voting record of government legislators.

What makes this all the more remarkable is that one of the original "hacktivists" is now Digital Minister of Taiwan – a relatively young trans woman at the time of her appointment in 2014. Audrey Tang brought into the state apparatus a concept that the problem of politics is one of *consensus*, and that this could be significantly improved by taking principles from open source software development. Tang identified a core problem with the internet at that time, which was that "the kinds of online spaces where political debate happened were engineered for an entirely different purpose: to capture attention" (C. Miller, 2020, para.8). The solution to this problem turned into vTaiwan⁴⁵, an online platform where government and people meet to define and discuss important issues.

The platform is designed to highlight where there is *consensus*, which is quite distinct from social media, for example, where people certainly can raise important issues but they tend to be overwhelmed by factional groups and derailed by algorithms built to maximise advertising revenue (Oremus, Alcantara, Merrill, Galocha, et al., 2021). The vTaiwan platform has already resulted in 21 cases resulting in decisive government action, including legislative reform. For example:

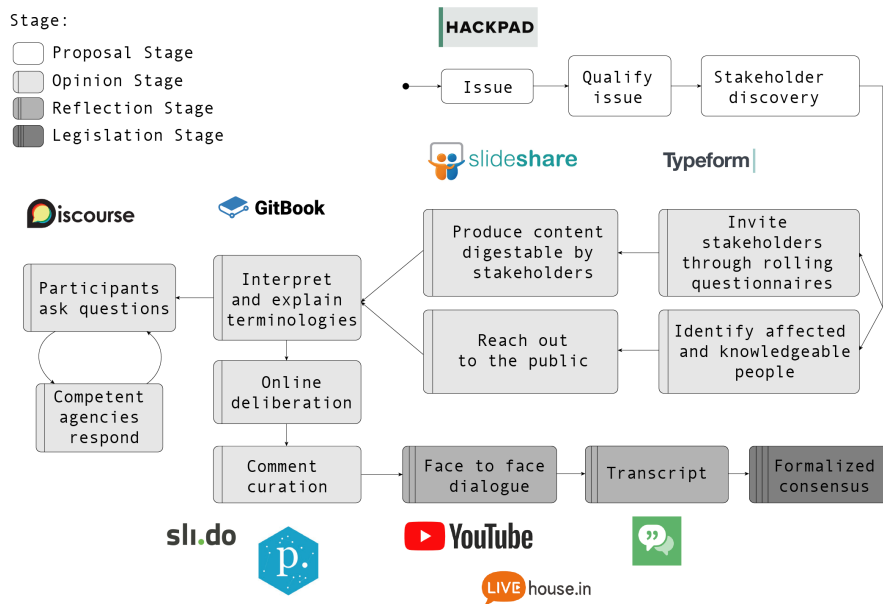
- **Uber regulation.** A crowd-sourced consultation led to regulatory reform of the taxi and ride-share sector in 2016. It was notable because, "instead of taking place behind closed doors ... it was livestreamed, and live-transcribed, with over 1,800 people watching. Faced

⁴⁴Refer to <http://budget.g0v.tw/budget> (accessed 7 May 2022).

⁴⁵<https://info.vtaiwan.tw/> (accessed 7 May 2022).

2. Literature review

Figure 2.4.: vTaiwan – how it works



with such clear public pressure, and knowing there was a real consensus behind the demands, Uber “caved in” on almost all of them” (Rashbrooke, 2021, para.9).

- **FinTech Sandbox Act.** vTaiwan was utilised to help draft and enact the Financial Technology Development and Innovative Experimentation Act 2018. This legislation provides a ‘sandbox’ for innovative new FinTech services to test, grow and, ultimately, bring their product to market.

A graphical representation of how vTaiwan works is shown in figure 2.4. In it we can see a range of online collaboration tools being brought together in a structured process to seek ideas, feedback and opinions. This is a formal process which results in *consensus*.

In terms of data sovereignty, Taiwan has embedded the principle of open data (albeit via a centralised system) and citizens are able to access *any* information that is held on them by government (Taiwan National Development Council, 2022).

In terms of legislation, the instrument most relevant to this thesis is the Personal Data Protection Act 2015 (PDPA)⁴⁶. The PDPA applies across government and non-government entities, although each have different expectations – with government having more discretion around collection and use of personal data. Some salient points, compared with other legislation we have

⁴⁶Official English translation is available at <https://law.moj.gov.tw/Eng/LawClass/LawAll.aspx?PCode=10050021> (accessed 7 May 2022).

2. Literature review

already looked at:

- There is no conception of ‘special data’, such as sexuality or religious belief, and Article 6 specifically exempts the processing of this information from PDPA oversight under a range of relatively loose conditions such as “where it is necessary for statistics gathering or academic research”.
- Consent for data collection can be presumed if the subject does not “indicate his/her objection”. Although the agency collecting data must still inform the subject around purpose and retention, for example, this requirement can also be waived in a range of scenarios, such as “where the collection of personal data is necessary for the government agency to perform its statutory duties” (Article 8).
- While there is a right to access personal data, there is also an exemption in the Act “where the material interests of the data collectors or any third parties may be adversely affected” (Article 10). This could be interpreted rather broadly, and appears to be relatively weak in the face of algorithmic fairness and related debates held in the GDPR context.

Although a very small selection, the above should indicate that there are a large number of broad exemptions in the PDPA and, indeed, work is ongoing to draft amendments that are more in line with the GDPR.

What we see in the case of Taiwan, therefore, is a unique approach to government transparency and consensus building, modelled largely on Free and Open Source Software (FOSS) principles. Legislatively, however, Taiwan is an entrenched centralised state (in terms of data management) and, while the achievements of g0v and vTaiwan are very impressive, there appears to be little discussion about data sovereignty. What is very much in Taiwan’s favour, however, is that a platform has been established upon which claims to data sovereignty could be brought; the question is whether or not this is an issue within the crowd-sourced consciousness.

2.2.5. The Aotearoa New Zealand context

Having briefly discussed data regulatory frameworks in other jurisdictions, we now turn our attention to Aotearoa New Zealand. Although it has ratified the International Covenant on Civil and Political Rights⁴⁷ (ICCPR), no formal absolute right to privacy exists in Aotearoa New Zealand law. It can (and indeed has) however been read in to the New Zealand Bill of Rights Act 1990 (NZBoRA), where several sections articulate freedom from such things as unreasonable searches and freedom of association (Penk, 2016b). This scenario is not uncommon in-

⁴⁷Article 17 of which provides that “no one shall be subjected to arbitrary or unlawful interference with his [sic] privacy”.

2. Literature review

ternationally; the US Constitution does not provide an explicit right to privacy, and many key international instruments (the ICCPR being one such example) also provide no absolute right. Competing interests must be balanced in a privacy claim, and hence formal language has tended to centre on “arbitrary” infringements of privacy.

Outside of the NZBoRA, the Privacy Act 2020 remains the key statute conferring privacy rights on New Zealanders. Although it is very much focused on privacy as applied to data and information, it arguably confers no strong protections in this specific area (Bartlett, 2021). The Act has numerous exceptions and is very welcoming in its understanding of ‘competing interests’. Furthermore it explicitly charges the Privacy Commissioner to weight those competing interests against claims, as well as the need for government *and business* to achieve their objectives. On a positive note, though, the 2020 amendment also included the need for the Privacy Commissioner to consider both “cultural perspectives” and international guidelines that may be “relevant to the better protection of individual privacy” (Privacy Act 2020, §21). This is a neat link back to our discussion on indigenous data sovereignty, although there are no examples available of the cultural perspective consideration being utilised thus far.

The original Privacy Act 1993 was ground-breaking when it was enacted, since it was one of the first laws globally to extend the scope of privacy outside of government and into the private sector. Penk (2016a) reflects on how it grew out of the 1960s and 1970s, where the overriding concern was the privacy impact, on private citizens, of information technology controlled by *governments*. Earlier legislation in North America reflected exactly this – the US Privacy Act 1974 and the Canada Federal Privacy Act 1983 were both concerned only with government. The Act today gives prominence to thirteen core principles, which are intended to encompass the full cycle of collection, storage and usage of personal data. These are worded rather generally and, as we have already noted, can be open to interpretation.

Whilst many of the exemptions referred to in the legislation are also a key part of other legislative frameworks reviewed in this section (for example, an exemption for use by security and intelligence agencies), their place in the Aotearoa New Zealand Privacy Act 1993 is heavily criticised by Penk who complains that they render the privacy principles “no more than ... ideals or aspirations, the application of which in any given fact situation may be uncertain, particularly as the exceptions may be interpreted subjectively” (2016a, p. 62). The exemptions were retained in the 2020 version and Bartlett concurs by noting that the Aotearoa New Zealand Privacy Act 2020 is “high level, without detailed or prescriptive rules” (2021, p.99) and is therefore quite distinct from the tight and prescriptive focus of the GDPR.

Actions brought under the Privacy Act 2020 can be remedied via a maximum financial penalty of NZ\$10,000 which pales in comparison to the sanctions available under the GDPR⁴⁸. Further-

⁴⁸Where Amazon and WhatsApp were fined €746M and €225M respectively in 2021.

2. Literature review

more, the Act is generally not enforceable in Courts of law – in the vast majority of cases, all complaints must be dealt with via the Privacy Commissioner.

Part 7 of the Privacy Act 2020 provides for the Privacy Commissioner to hold an oversight role across formal government ‘data matching’ programmes, a function that has been carried over from the previous version of legislation. The Commissioner plays no part in agreeing or assessing new data matching initiatives but simply reports on them and, under §158 of the Act, can advise on the justification for a programme’s existence – 55 such programmes are active at the time of writing ⁴⁹. On the topic of data matching Penk (2016a) notes that, despite some clear advantages (such as fraud detection, or maintaining currency of data) such data matching can simply perpetuate poor quality data and – from a data sovereignty perspective – individuals are unable to correct or control their information, or understand exactly how it has been utilised. Furthermore whilst the Aotearoa New Zealand Privacy Act 1993 was originally praised for expanding the scope of privacy into the private sector, the Privacy Commissioner holds no jurisdiction whatsoever over data matching arrangements that are not occurring solely between government agencies.

The actual protections conferred on individuals may be diluted by this part of the Act. It is probable that data matching explicitly involves the use of information for a purpose *different than that for which it was collected*, and thereby contravenes Principle 10 (Limits on use of personal information). Furthermore, the information has not been collected directly from the individual and thus would also contravene Principle 2 if exemptions for “avoiding prejudice” to law enforcement and the “protection of public revenue” were not also in place (Privacy Act 2020, §22). The co-existence of these tensions is axiomatic in the context of the Act, where the Privacy Commissioner is required to balance a range of competing interests – largely due to its “high level” nature (Bartlett, 2021). Where such data matching programmes are implemented, and Penk’s concerns (noted in the preceding paragraph) are realised, we must assume that the Privacy Commissioner has determined the right of government to achieve their objectives has trumped the protections granted to individuals and, in fact, this is strongly indicated in the wording of the available exemptions.

The 2020 amendment did introduce some new features which are relevant to our discussion about data sovereignty. A recognition of the growing momentum behind cloud computing prompted the inclusion of a new Principle 12 – “Disclosure of personal information outside New Zealand”. This provides for the movement of personal information overseas, only if the transferring agency believes that the target jurisdiction will “provide comparable safeguards to those in this Act” (Privacy Act 2020, §22, IPP 12(1)(c)). This is an extremely ‘light touch’ approach

⁴⁹Please refer to <https://privacy.org.nz/privacy-act-2020/information-sharing/information-matching-provisions/> (accessed 5 May 2022).

to data localisation, which clearly signals that the government do not see an existential geopolitical data sovereignty threat, but are primarily concerned that the personal information of New Zealanders will be adequately protected no matter where it resides.

2.3. Centralisation, decentralisation and distribution

“Centralized stores of data are a surveillance state’s dream” (Raval, 2016, p.15).

In this section we will turn our attention to the specific forms of centralisation, decentralisation and distribution, and the implications of each as it relates to this research. We will delve further into what the literature has to say on this rapidly developing topic, before surveying the landscape of technologies and frameworks that offer an alternative to centralisation. Finally, I will assess candidate technologies and determine which of these is to be utilised for the prototype artifact.

2.3.1. Understanding network types

Whilst we have introduced some key concepts already in section §2.1, they are central to this thesis and it is worth exploring them further for two reasons:

- To ensure we have a clear understanding of these central concepts before progressing in more detail
- Because terms and concepts are used interchangeably (and in my view incorrectly) in the literature, and meaning can be inconsistent.


One of the seminal thinkers around these topics is Paul Baran, who worked at RAND Corporation in the USA from 1959, and invented one of the bedrocks of computer network communications (and the internet) that is still used today⁵⁰. His 1964 paper is one of the first to identify and present a taxonomy for different kinds of computer networks, as shown in figure 2.5.

When designing this representation Baran, of course, was not thinking about all the technological advances available to us today. The primary concern was to understand how to ensure the security and reliability of a given communications network. In this way, he is entirely right to state that “since destruction of a small number of nodes in a decentralised network can destroy communications, the properties, problems and hopes of building a ‘distributed’ communications network are of paramount interest” (Baran, 1964, p.3). The aim at that time was to spread risk, and limit points of failure. Baran considered that both the centralised and decentralised models,

⁵⁰This is the notion of ‘packet switching’, or the grouping of data together into digital packets for transmission.

2. Literature review

Figure 2.5.: Centralised, decentralised and distributed networks (Baran, 1964)

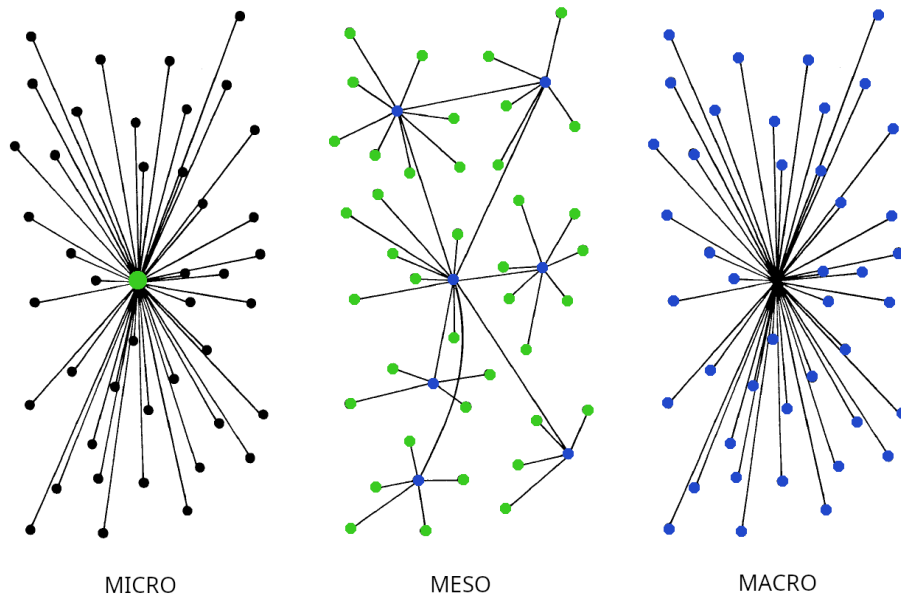


as represented in figure 2.5, have inherent weaknesses in this regard. That is, the termination of any single central node would render connected nodes obsolete. Only in the distributed model does a network have the resilience to cope with the failure of multiple nodes. Baran was writing about *network security* and we can therefore see that one's frame of reference appear to be highly relevant to understanding. Hoffman et al. (2020) expand on this and note that there cannot be a single universal conceptualisation – context is everything and domain-specific research is vital. I also argue that the *level of analysis* is just as important.

By this I mean that in any moderately complex domain, there will be a range of layers that should be analysed separately; this thesis is concerned with the healthcare domain. If we imagine a single primary care clinic it is easy to think about Baran's model A, whereby the central node is a Patient Management System (PMS) and outlying nodes are patients or other health providers who are submitting or transferring information. The same model broadly applies to a hospital PMS. So, at the most detailed layer of analysis available in healthcare (a patient visiting a health facility to receive some service), things appear to be highly *centralised*. This is shown in figure 2.6 as the 'micro' layer. If we move up a level, and think about the broader health sector, we may consider that Baran's model B is more applicable. At the 'meso' layer we can see the green health facility nodes converging their data in DHB or regional systems, shown in blue. DHBs themselves have some connectivity to share information, so that patients can be treated regionally across secondary care and tertiary specialist services where necessary. If we zoom

2. Literature review

Figure 2.6.: Health sector layers, adapted from Baran (1964).



right out, then the health data ecosystem will look more like model A again – where data flows to the centre. This is shown as the ‘macro’ layer, where DHBs and regional systems aggregate data to send to the centre, which is represented by what appears to be a black hole.

So what we tend to see in practice is differing configurations of centralised data stores, sometimes connecting to others, but with a general trend for data to flow in centrally to government agencies. Baran’s representation of ‘decentralisation’ in this case gives too much credit to the *meso* layer where, in practice, the data remains highly centralised. But we should once again remind ourselves that context is everything, and Baran certainly had a different focus than data sovereignty.

It may be interesting to note at this point that the original design of the modern internet’s first iteration – the ARPANET (funded and established by the US Department of Defence in 1969) – was in fact explicitly about connectivity, redundancy and *decentralisation*. The project director Charles Herzfeld later noted that “the ARPANET came out of our frustration that there were only a limited number of large, powerful research computers in the country, and that many research investigators, who should have access to them, were geographically separated from them” (quoted in Gallagher, 2019, para.2). If we are viewing the internet as a knowledge commons, therefore, we may see its early iteration as being decentralised in terms of *access* to information, but centralised in terms of *ownership and control* of the information. There were a limited number of computers which hosted information, but the prevailing model was that these should

2. Literature review

be accessible to researchers. Similarly, in healthcare, we see a large number of connected data stores, where the prevailing model is for data to move towards the centre. Zuboff (2019) has written about how the spirit of the internet has shifted in the modern era, following recognition of the power of data and its value.

This introduces another layer to the analysis, away from a simple view of the *technical* architecture, by introducing a concept of the *philosophy* that sits behind a model. That is, what is our philosophical approach towards ownership of data and how it can be accessed? We may have a philosophical commitment to individual access and ownership of data, but are technically stymied by the tools and systems in play. This appears to be where MOH are at currently, as they embark on the *Hira* programme which aims to transform sector platforms to enable a decentralising of data (which I contend does still not fully meet the spirit of data sovereignty, but is a very positive step in that direction).

To help us locate this debate within a broader theoretical context, we can use the important work of Ostrom around common-pool resources and managing commons (Ostrom, 2015). In this book, we see a focus on a third way – between Leviathan or privatisation – in a discussion around managing finite collective resources. This third way can be summarised as “self-governance by resource users” (Herzberg, 2020, p.629) and places great emphasis on distributed and localised solutions. Ostrom’s concern around the need for self-governance, and individual freedom, aligns very well with the principles of data sovereignty⁵¹.

In chapter 1 I have set out the notion that greater access to data is likely to engage more people in their healthcare, and has the potential to produce more positive health outcomes than a centralised government problem-solving approach – particularly in groups that are already marginalised from mainstream policy analysis and debate. Taking a slightly different angle, Levine (2011) uses the example of a university based *associational* knowledge commons as a means to engage students and citizens in work of public value – the project had the benefit of exposing people to important public issues, at the same time as increasing their engagement with them and the commons itself.

Whilst Ostrom wrote as a political economist, mainly concerned with natural resources, other writers have expanded on her thinking and brought these principles into very diverse arenas – including data and, specifically, health data. In fact Ostrom’s thinking becomes more relevant in the information or knowledge arena, since these things are non-rivalrous and, when access is opened, they become a universal public good. In any case, even where goods are not non-rivalrous, “groups can effectively manage and sustain common resources if they have suitable conditions, such as appropriate rules, good conflict-resolution mechanisms, and well-defined

⁵¹I would include indigenous data sovereignty here, save for Ostrom’s focus on *individual* freedoms which may clash with indigenous epistemology regarding the collective.

2. Literature review

group boundaries” (Hess & Ostrom, 2011, p.11). In discussing a Medical Information Commons (MIC), Bollinger et al. (2019) reached several key conclusions:

- Data for a MIC should not be centralised, and its operation would depend on linking disparate data
- Agent-centricity was key to success, reflecting the notion that governance of a common-pool resource should be the result of collective agreement by stakeholders (Ostrom, 2015)
- Data scope should go well beyond the very basic administrative data sets commonly used by governments, to include patient-generated health data
- Insights or outcomes derived from the MIC must be contributed back into it.

Applying these principles to a national health information system does indeed feel radical, and yet these conclusions are achievable – *only* achievable, in fact – within a distributed paradigm. With the best will in the world, an ecosystem effectively owned and controlled by even the most benevolent entity will inhibit the self-determination and self-organising principles that Bollinger et al. (2019) advocate.


2.3.1.1. The centralisation continuum

The foregoing discussion highlighted that understanding of concepts, and their application in practice, is heavily mediated by context, unit of analysis and also distinguished between technical and philosophical approaches. Hoffman et al. add to this complexity by suggesting that centralisation-decentralisation is a false dichotomy – “the binary positioning of ‘decentralization’ and ‘centralization’ is not an ideal approach as it obscures the motivations and rationales for the process” (2020, p.15). Their focus is on blockchain, and the extent to which it may be considered decentralised, and conclude that any blockchain will have characteristics of each end of the continuum. They conclude by noting that decentralisation may not simply be the opposite of centralisation, and question “whether the two processes are qualitatively different in ways that would disallow treating them as inverses of each other” (Hoffman et al., 2020, p.15). One of the most prominent ‘decentralised’ technologies today is blockchain, so it is instructive to explore this further.

In the case of the ‘decentralised’ architecture (item B in figure 2.5), we can see that groups of nodes are able to connect with a single central group node – but not directly with every other node. This is a good representation of a blockchain network, where each group node (shown coloured red in figure 2.7) hosts a copy of the ledger and users connect to that node to interact with the blockchain. It is possible for individual nodes to connect remotely – for example the

2. Literature review

Figure 2.7.: Blockchain as a decentralised network



node at the very top is able to communicate with the node at the very bottom. However, this communication must be routed via three group nodes. In Baran's taxonomy, the failure of any of these will isolate connected individual nodes and break that communication. We can therefore say that there exists a type of technical centralisation within blockchain, by virtue of a users' reliance on connecting to a node – a server which is under the ownership and control of an entity.

Hoffman et al. (2020) consider that 'decentralisation' is about moving power away from a central authority and, in this sense, blockchain talks a very good game. We reviewed in section 2.1.4 the prevalence of the 'decentralisation' terminology in blockchain marketing. Following Baran's analysis I contend that, while blockchain does move information away from a single source, it simply *replicates* it across other central nodes. As we can see in figure 2.7, rather than one central node there are now seven centralised nodes, each with an exact copy of the same data. This has advantages in a cryptocurrency domain (as we discuss in section 2.3.3.1), but offers us relatively little advantage when we wish to realise health data sovereignty.

Two examples provoke us to question how decentralised cryptocurrency truly is. Firstly, the Ethereum cryptocurrency suffered from exploitation of a smart contract in 2016, where US\$50 million was effectively stolen. As a response, Ethereum developers created a replica of the existing blockchain but transferred the stolen funds to a recovery address thus 'recovering' the

2. Literature review

money (Buterin, 2016). Similarly, in May 2022, investors in the Juno cryptocurrency determined that US\$36 million had been acquired by another investor maliciously and a decision was made to remove this money from that user’s wallet. Unfortunately the developers carrying this out made an error and accidentally transferred all the currency to an inaccessible address (Boom, 2022). What is striking about these two examples, is exactly how much centralised control the developers have over a cryptocurrency. Certainly, it may be decentralised in the sense that it exists and operates entirely outside the state apparatus. But it is confusing that so much emphasis is placed on the language of ‘decentralisation’ when a user’s money can be moved around at will by a developer.

Complicating matters still further, we have Baran’s notion of a ‘distributed’ network which was the ideal as far as his research was concerned. In Baran’s taxonomy, every node is theoretically able to connect with every other node (even if only indirectly). In modern applications, ‘distribution’ tends to mean that the whole data set is distributed amongst a group of peers – giving the appearance of a single whole (whether it is serving files or performing computational tasks), but actually being spread across many⁵². We can take a lead from the ‘distributed computing’ approach to understand this more, where it has been defined as “a collection of autonomous computing entities connected to accomplish a joint task” (Perrin, 2017, p.12). Within this understanding of ‘distribution’ it is possible to understand the entirety of a domain, but by default no individual nodes can do so. If we compare this to ‘decentralisation’, then it is clearly much less about connection because of the in-built bottlenecks represented by the group nodes. Hoffman et al., in fact, demonstrate some confusion about the application of these terms against technical architectures, claiming that “a chiefly technological classification would categorize blockchains as a type of peer-to-peer and a serverless distributed system” (2020, p.2) which, I respectfully suggest, is almost completely backwards. However, they rightly go on to note that understanding must go above and beyond technical architecture, which we have already touched upon by introducing the concept of *philosophical* centralisation.

This all reiterates the point that our conceptual understanding depends on our frame of reference, and the layer of the issue we are addressing. This is summarised in table 2.4, where I put the three identified paradigms together with characteristics across both technical and philosophical dimensions. What, I hope, this makes clear is the following key points:

- A network is not truly distributed in any form, as long as any entity is making ownership or control decisions on behalf of others
- Self-determination and self-governance can only be enabled with integrity by a distributed model

⁵²We will use the example of BitTorrent in more detail in section 2.3.2.3.

2. Literature review

Table 2.4.: Centralisation matrix

Paradigm	Technical characteristic	Philosophical characteristic
Centralised	One central node holds all the data	A single entity controls access to the data
Decentralised	Multiple central nodes hold overlapping or replicated data	Each node controls access to its data
Distributed	All data is distributed across all peers who connect directly with each other	Access and ownership is self-determined

- Whilst philosophical and technical characteristics can be mixed, they can only do so to a limited extent. For example, any philosophical intention to permit decision-making authority over data will always hit a hard limit within a centralised paradigm.

Hoffman et al. (2020) note that literature dealing with blockchain and decentralisation tends to focus on the practicalities of *how* it is achieved, and the perceived benefits. Understanding of decentralisation in this context is furthermore focused on what can be better described as ‘disintermediation’. That is, building a network that has traditionally been centrally managed by some type of state apparatus but, via use of blockchain, enabling this governance layer to be removed. But Hoffman et al. (2020) also caution that this does not necessarily guarantee decentralisation where we now see a growth of intermediary organisations operating as *de facto* gatekeepers.

We have already noted the highly centralised model behind the NFT marketplace, but we should also recall the enormous power that blockchain developers appear to have in determining – for all blockchain’s marketing about immutability – where currency should reside (Boom, 2022; Buterin, 2016). We are therefore left wondering what the real goal of cryptocurrency is. If users are indeed happy to invest in a marketplace which sits outside the state apparatus, but is in many cases highly centralised and subject to direct intervention by developers, we have to conclude that the overriding goal is simply the removal of state apparatus such as a central bank (H. Wang et al., 2018). Walch (2019) sums this up neatly by referring to a “veil of decentralization” in the following way:

“Despite the common use of “decentralized” to indicate that power is diffuse rather than concentrated in a blockchain system, existing blockchains ... have small co-ordinated groups who shape how the systems operate... there are many parts of blockchain systems that are exceedingly centralized” (Walch, 2019, p.62).

In summary, concept definition can be complex and there seems to be varying levels of understanding, particularly regarding blockchain, around what constitutes useful examples of each.

2. Literature review

This research is aiming to design a health information system with no centralised oversight and, therefore, we are seeking a truly distributed solution.

2.3.2. Why is distribution important?

Having established some of these key concepts, we must now answer the question – why should I care about any of this? The reality, as we know, is that many people either believe they have better privacy protections in place than they actually do (Hoofnagle & King, 2008), or they hold fatalistic or ambivalent attitudes about use of their personal data (McMullan, 2015; Solon, 2018). Furthermore, we know that the centralised paradigm is dominant and – based on selected measures – generally operates very effectively. Whilst centralisation is dominant, the fact is that distributed approaches to computing have a venerable history and have led to the creation of tools and products which are still core to the current technology landscape today. In this section, I will review a selection of these in order to establish the practical relevance of distributed approaches.

2.3.2.1. Free as in speech; free as in beer

1983 saw the birth of the FOSS community, when Richard Stallman – an Artificial Intelligence engineer at Massachusetts Institute of Technology (MIT) – unveiled the ‘GNU’ project. The primary goal of the project was to bring a completely free operating system into existence for the very first time.

‘Free’ was to be considered both in monetary terms (cost to the end user), as well as philosophical terms, which manifested as the ability for people to share the code, and modify and publish their own versions for others to use. The ‘free speech’ component was realised in 1989 with the release of the GNU General Public Licence (GPL) which explicitly targeted two of the key methods by which restrictions were being imposed on software at that time. Firstly, it insisted that software is published in human readable form (so it can be inspected for vulnerabilities, and utilised for learning) and, secondly, it ensured that licenced software could only be used with other software when it is licenced under a more permissive model. Or, in other words, a more restrictive licence could not ‘trump’ the GPL, were it to be combined with open source software⁵³. But what does all this have to do with the distributed paradigm?

Following on from the release of the GPL, communities and projects were brought to life centred around this notion of transparency and visibility based on philosophical freedom. A shift

⁵³Richard Stallman has since become a vociferous critic of ‘Open Source’ since, in his view, it moves too far from FOSS principles and does not explicitly strive for free (as in beer) software – “We libre-software activists say, ‘Software you can’t change and share is unjust, so let’s escape to our free replacement.’ Open source says only, ‘If you let users change your code, they might fix bugs.’ What it does say is not wrong, but weak; it avoids saying the deeper point” (quoted in Vaughan-Nicholls, 2018, para.14).

2. Literature review

occurred from a ‘top down’ model of software development – where work was done effectively in secret by large corporations and released to the public in non-readable binary formats – to a ‘bottom up’ model, where new operating systems and software were being developed remotely by volunteers who published all of their work for anyone to access and use. Eventually this developed into the ‘Linux’ operating system, first released in 1991, and which today powers two thirds of all websites (W3Techs, 2022), Google’s ChromeBooks, the Android operating system, and a majority of routers, modems and other embedded devices. The point of this is to state that many core technologies that are now relied upon were only brought about by the move to distribute control and direction of software development. The FOSS example shows us that both the technical *and* philosophical components are important for the distributed model to become really effective.

Many writers have expounded on the impact of FOSS, often from a strong political or philosophical position. This is most eloquently realised in Eric Raymond’s 1999 book ‘The Cathedral and the Bazaar’, where the author argues that ‘bottom up’ (the Bazaar concept) provides huge benefits across a range of domains when compared with the ‘top down’ (the Cathedral concept) approach (Raymond, 2001). This essentially boils down to some very practical points around software development – not least the fact that, the more people you have reviewing the code, the more likely you are to find and fix problems. Furthermore, however, Raymond makes other relevant points like “if you have the right attitude, interesting problems will find you” which, in the FOSS context, is a long way from the Cathedral model, driven by corporate interests and shareholder returns. The fact that this approach has turned into something so ubiquitous, and something impacting so many people, is testament to its efficacy. Affirming this, Schweik has noted that “the collaborative ideals and principles ... [of FOSS] ... could be applied to any collaboration ... and could potentially increase the speed at which innovations and new discoveries are made” (2011, p.277). If we can reconcile the core FOSS principles with distribution, then we are saying here that distribution itself has the ability to effect meaningful change within society.

2.3.2.2. Git

Git is a distributed version-control system, used to coordinate software development and track changes to files. It was created in 2005 by the creator of Linux, Linus Torvalds, after the cost-free status of the proprietary version-control software being used at the time was revoked. This presented an opportunity to use key lessons learned up to that point, but Chacon and Straub describe its creation thus – “as with many great things in life, Git began with a bit of creative destruction and fiery controversy” (2020, p.13). It also reflects Raymond’s notion that, given the right attitude, “interesting problems will find you” (Raymond, 2001).

When working in software development, often a team of people need to be interacting with

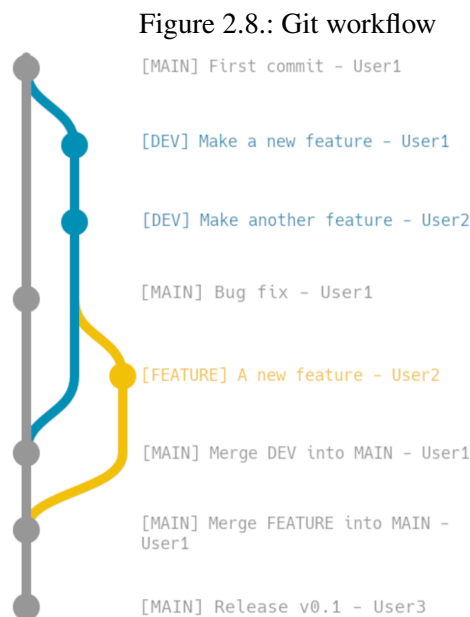
2. Literature review

different parts of the code base for testing, in such a way that the existing ‘live’ version of the product is unaffected. Version control also allows you to revert back to any prior state, as well as providing clarity about exactly who did what and when (Chacon & Straub, 2020). An overview of a Git workflow is shown in figure 2.8.

In this example, some software is being built. Working from top to bottom, User1 creates an empty repository and places the initial code base in a branch called MAIN. The same user then saves that code in a separate parallel branch called DEV. This allows all the code in MAIN to continue unaffected, while some additional work is tested. We can see that both User1 and User2 then build some new features in this DEV branch. While they are testing those, User1 notices a bug in the existing code in MAIN and fixes it. User2 then takes the code in the DEV branch, and saves it in another separate parallel branch called FEATURE. User2 builds yet another new feature there. Before testing on that is finished, User1 merges the changes in DEV back into the MAIN branch. The MAIN code base, now has the two improvements built by User1 and User2. Subsequently, the new feature in User2’s FEATURE branch is merged directly back into MAIN. Finally, the MAIN code is released as an official v0.1 by User3.

The creation of new branches (as seen in above with DEV and FEATURE) is the true distributed power of git. Anyone can take existing code of a working application, ‘fork’ it to another branch, make any changes they want and then – as long as it is a genuine improvement that fits with the aims of the original project – ask for it to be merged back into the main code base. There is another important scenario where developers from a project disagree about the direction, and a ‘forked’ project ends up becoming something completely different but with its origins in the original shared code base. Some notable examples of this are:

- 2009 – MariaDB is forked from the MySQL⁵⁴ code base, due to the acquisition of MySQL by Sun Microsystems and community concern around their plans for it.



⁵⁴A very popular database engine used in many websites.

2. Literature review

- 2010 – LibreOffice⁵⁵ is forked from OpenOffice.org as a result of perceived neglect of the code by Oracle.

Git is free in both dimensions advocated by the FOSS community, and the full code base for it is available at <https://github.com/git/git>. Many different web-based services have evolved, acting as a front-end for git and also providing storage space for code repositories. One of the most prominent – GitHub – claims it holds more than 200,000,000 code repositories, and is used by more than 4,000,000 organisations and businesses globally⁵⁶. GitHub was acquired by Microsoft in 2018 for US\$7.5 billion⁵⁷ (Microsoft, 2018).

One of its core original purposes was to be a “fully distributed” system (Chacon & Straub, 2020), meaning that all data and workloads could be shared across any number of nodes or users. In the case of Git, I can download any code repository I want, and start to make changes to it and share it with other people; the knowledge and information has been distributed freely. Hemel and Coughlan (2017) contend that Git’s ability to freely distribute code and content qualifies it as being *completely distributed*.

While it is true that only one Git repository at a time is considered authoritative this may simply be a matter of perspective. In practice, there will only ever be one ‘Linux’ repository and accepting new code into that project is famously difficult (McMillan, 2012). But this doesn’t necessarily minimise the distributive power of git; even if I fork code and my changes are not accepted into the authoritative live repository, I still have my own code that reflects my hopes and aspirations and I can expand and share it as I please.

One issue to note is the growth of cloud intermediaries, such as GitHub and GitLab, which are built on top of Git and on which many people (as we have seen) rely. Whilst it is entirely feasible to set up and run your own version of Git, these intermediary services are appealing by virtue of being free (up to a certain threshold) and because retention and storage of your codebase is outsourced. There are two current issues here. Firstly, in 2019, GitHub confirmed it was blocking access to users based in Iran, Crimea, Cuba, North Korea and Syria due to being subject to US trade control legislation (Porter, 2019). GitHub has also blocked access to specific pages for users from certain countries, under pressure from their governments (Lunden, 2014). Secondly, in 2021 GitHub released a new service called ‘Copilot’ which offered the ability to have your code automatically generated via AI. As with any AI, the output merely reflects the

⁵⁵A FOSS replacement for Microsoft Office.

⁵⁶As per <https://github.com> (accessed 7 May 2022).

⁵⁷Interestingly, announcement of the acquisition was followed by a large exodus of users from Github to their close rival Gitlab, in protest at Microsoft’s perceived move to control the Open Source community (Sharma & Mukherjee, 2018). Former Microsoft CEO Steve Ballmer is on record describing Linux (and the GPL) as “a cancer”, and the company felt sufficiently moved to issue a public apology in 2020 (T. Warren, 2020).

2. Literature review

training data that has been used as an input. In this case, Microsoft had utilised all the existing code held on its servers – without seeking consent from users – as training data. This introduces three separate issues:

1. Licensing. Code on GitHub can be assigned a licence. Copilot does not indicate the licence of the code that was used to train a response, leading to concerns about licence violation (Gershgorn, 2021).
2. Efficacy. Pearce et al. (2021) discovered that approximately 40% of Copilot suggestions contained security vulnerabilities.
3. Privacy. Reports surfaced of Copilot suggesting private information, such as phone numbers and Application Programming Interface (API) keys, to developers (Anderson & Quach, 2021).

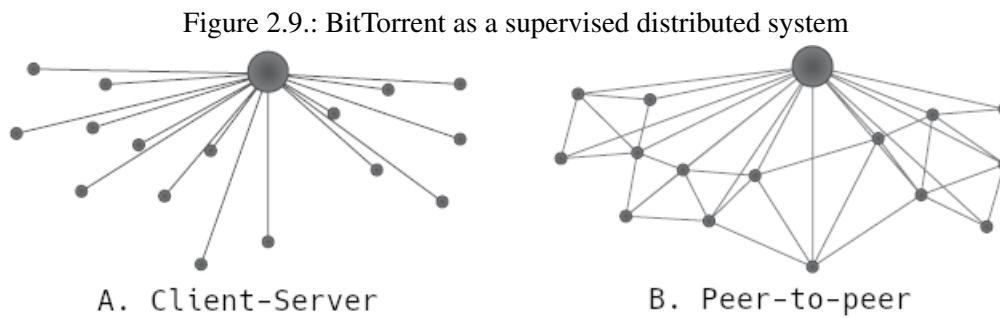
2.3.2.3. BitTorrent

BitTorrent is a protocol for file-sharing, made available in 2005, which transcends the limitations of a traditional client-server approach by directly connecting all peers together.

In a traditional approach, there may be one single web server which hosts files for people to download. If a large number of people connect directly to that one machine to download the file then download speeds will be adversely affected or, worse, the download service may fail completely. This approach defines a single centralised point of failure.

BitTorrent, by contrast, solves this problem by breaking files into chunks which are shared randomly with anyone wanting to download a file. If a file has 100 chunks, for example, then Downloader 1 may be given chunks 1-50, and Downloader 2 may be given chunks 51-100. At this point, the entire file has been distributed to the downloaders but none of them have the whole file. To reduce burden on the web server, Downloader 1 can now simply get the remaining chunks (51-100) *directly* from Downloader 2, and vice versa. The power of this approach becomes exponentially more beneficial where a large number of users (in this configuration referred to as a ‘swarm’) can offer multiple copies of the same chunk to new users, thus increasing the speed significantly. The software used by participants of a swarm keeps track of exactly which chunks it has, and which it still needs. In this way, it is possible to stop a transfer halfway through and resume months later without any loss of data integrity.

BitTorrent is a FOSS project and the full source code is available at <https://github.com/kenorb-contrib/BitTorrent>. While the protocol has been controversial due to its association with illegal download of copyrighted content, it should also be noted it is simply a protocol for file sharing and it has many legitimate uses. Most Linux distributions, for example, offer



to download the operating system⁵⁸ via BitTorrent. Similarly, <https://academictorrents.com/> is a service designed to leverage BitTorrent to distribute large data sets (totalling more than 127TB) for Science researchers.

The notion that users in a BitTorrent ‘swarm’ connect directly to each other suggests that it uses a distributed model, although some method of connecting peers together is required. To achieve this, all users connect initially to a central coordinator (a web service referred to as a ‘tracker’) which passes back to each client a directory of other users. From this point, users can communicate directly with each other but only where the tracker is monitoring and passing information back to peers about which users have which chunks. This results in the ability to conduct ‘supervised’ peer-to-peer communication, as shown in figure 2.9.

The very low latency and rapid validation of BitTorrent are clearly of interest when thinking about building a scalable distributed information ecosystem. BitTorrent represents an impressive example of a FOSS project which rejects centralisation, and has set a precedent for the feasibility of peer-to-peer communication at scale.

2.3.3. Candidate distributed approaches

The previous section outlined that there is a very meaningful precedent for the distributed paradigm and that those technologies are, in some cases, a key part of modern computing. The aim was to show that many core concepts have been proven. A key aim of this thesis is to use a modern framework which offers the ability to build distributed apps, so that we can test the feasibility of a distributed health information system.

In this section I will briefly review some noteworthy technologies that utilise fully distributed or decentralised approaches, and might offer the potential to build a radically distributed health information system. These technologies have been sourced via review of Twitter Bluesky’s

⁵⁸Usually relatively large files in the order of ~1.2GB. One prominent example can be found at <https://ubuntu.com/download/alternative-downloads>.

2. Literature review

‘Ecosystem Review’ (Graber, 2021), a general review of the literature and my own awareness of the current technological landscape.

2.3.3.1. Blockchain

Firstly, it is important to point out that blockchain is really a ‘methodology’ rather than a specific product or technology and one that has received a great deal of research interest⁵⁹. There are many different variations on the concept, and it has been applied in many differing ways. The key point to note, in order to understand the core premise of blockchain, is that of the ‘distributed ledger’. While all blockchains are distributed ledgers, not all distributed ledgers must use blockchain; it is merely a specific way to implement a distributed ledger.

The point of distinction for blockchain is in the way it records and stores entries into the ledger. This is done by cryptographically signing every transaction, so that a majority of peers in a network can verify that the transaction is valid and, simultaneously, that no previous ledger entries have been tampered with (Li, He, & Haiquan, 2021). Doing this in practice requires a large amount of overhead and, in the case of cryptocurrency, very large amounts of computing power to complete a validation and be rewarded with a small amount of currency (a process known as ‘mining’). This notion of consensus is critical to blockchain and, understandably, is vital where a large number of peers in a network are exchanging money or other items of value. In the domain of a ‘permission-less’ network (one where anyone can take part, and anyone can view the entire ledger), blockchain really excels with its focus on validating transactions, and ensuring that consensus means a group of bad actors cannot easily hijack a network. It is easy to imagine these being critical features of a cryptocurrency implementation. Health data, though, has very different requirements.

While we certainly do not want records to be tampered with, the big difference is that a health information blockchain would never be ‘permission-less’. If health information was to be stored directly on the blockchain, it would need to be designed in a way that a majority of the content was not accessible to most people. Furthermore, we can imagine that entities who would be writing data to a healthcare blockchain (your doctor, a hospital care team, or paramedics) would already be considered ‘trusted’ entities and therefore validation of their transactions becomes mostly redundant (as long as our authentication process is robust). Thinking through a healthcare example like this, I would suggest that two of the key features of blockchain become irrelevant in the healthcare context:

- **Transaction validation.** We will already trust healthcare providers to write accurate data to our health record. Opening our health record up to third parties, such as Apple Health

⁵⁹The AUT Library database shows a total of 29,604 books and articles published on the topic of blockchain in 2021.

2. Literature review

for example, would change this picture considerably but, for now, this feature offers no advantage if we simply want to reshape ownership and control of our health information.

- **Consensus.** In a large network of peers, consensus would be vital. This feature becomes unwieldy, however, when we consider that most individuals would be having a relatively small number of entities in their own health ‘network’. For example, I may only ever share my data with my General Practitioner (GP) and a hospital outpatient department if I have a chronic condition. Implementing a fully-fledged consensus process across a bilateral data sharing arrangement between my GP and I is not required, and makes the implementation needlessly complex and inefficient.

We have already outlined the doubt of researchers such as Walch (2019), in terms of where blockchain really sits on the centralisation continuum. In order to flesh this out a little more, we can look more closely at how a blockchain network operates. Any blockchain network requires the following components⁶⁰:


- **Membership Service Provider (MSP).** This is a service that enables management of identity and verification/authorisation in a network. The MSP must be able to issue and sign certificates verifying the identity of users and their transactions. This service must run on a server of some kind for an organisation.
- **Ordering service.** In permissioned blockchain networks, a separate service must run which guarantees deterministic consensus around the validation of blocks. Some kinds of blockchain (Bitcoin or Ethereum, for example) are not permissioned and do not require a separate ordering service. If one is required, it must run on a server of some kind for a network where there may be several Orderer nodes.
- **Peer nodes.** These are really the core of a blockchain network. Each peer node will host a full and complete copy of the entire blockchain; the ‘consensus’ is determined between any peer nodes within a network. Peers also host code which determine exactly what end users are allowed to do.

Each of the three entities above need to run on a single server instance; a single blockchain network may feature any number of each of these. End users who may want to get data from a blockchain, or update it, are required to connect to a Peer node. The entity that owns the peer node they connect to potentially has full visibility of any data in the blockchain although, in

⁶⁰Based on the Hyperledger Fabric blockchain implementation model (<https://hyperledger-fabric.readthedocs.io/en/latest/>), accessed 7 May 2022), but mostly applying to different permissioned blockchains too.

2. Literature review

Figure 2.10.: Example Hyperledger Fabric blockchain network



practice, the data would be encrypted or, in the case of Hyperledger Fabric, it would be permissioned. An example representation of a Hyperledger Fabric blockchain is shown in figure 2.10.

In this diagram, we can see the following items which represent the need for a centralised server of some kind:

- CA1-4. These are the four MSPs in this network.
- P1-3. There are three separate Peer nodes in this network.
- O4. There is one Ordering node in this network.

The end users (A1-3, shown to the left) each have to go via a Peer node and the Ordering service to interact with this blockchain. It is certainly true that the blockchain itself does not depend on a centralised authority (Geekyanage Don & Motalebi, 2021; Peterson, Deeduvanu, Kanjamala, & Mayo, 2016), yet we can see from the above that in practice end users are dependent on intermediary authorities (the entities managing Peer nodes) which hold distributed copies of the ledger. The extent to which blockchain is centralised or decentralised is largely a product of the type of blockchain chosen, and its scale – but it is never *distributed*.

To help us think about this, H. Wang et al. (2018) consider three types of blockchain: Public, Private and Consortium. Examples of Public blockchains are the Bitcoin and Ethereum cryptocurrencies; anyone can interact and use these products. These are thought to have resilience and redundancy, precisely because there is a large amount of users interacting with the network – or, using a FOSS analogy, “given enough eyeballs, all bugs are shallow” Raymond (2001, p.30). Private and Consortium blockchains overlap with each other somewhat, but the key point

2. Literature review

is that they are not open to any number of users by default and, thus, control and management of the network has been distilled into fewer hands. If we imagine a blockchain which underpins a RDHIS, it is easy to see that there would likely be relatively few key actors in the network (certainly compared with an example such as Bitcoin) who would be operating peer nodes, and be able to validate transactions. A summary of these three blockchain types is shown below:

Table 2.5.: Comparison of blockchain types. Adapted from H. Wang et al. (2018).

Property	Public	Consortium	Private
Determining consensus	All miners	Selected nodes	One organisation
Read permission	Public	Public / restricted	Restricted / may be public
Immutability can be compromised?	Nearly impossible	Possible	Possible
Centralised	Low	Partial	High

We can see here that distribution in a non-public blockchain network is likely impossible. In fact de Aguiar et al. identify a general “tendency to centralisation” (2020, p.369) in blockchain as something that must be addressed. Notwithstanding this, there has been great general interest in assessing blockchain for use in the healthcare sector, mostly as a means of sharing data between centralised authorities.

Peterson et al. (2016), for example, have suggested storing patient data in a blockchain but – rather than storing the actual data on the blockchain itself – writing only a FHIR URL (effectively a unique pointer to a piece of information, held by a centralised entity such as a hospital). This has the perceived merit that it “allows institutions to retain operational control of their data, but more importantly, keeps sensitive patient data out of the blockchain” (Peterson et al., 2016, p.3). This offers us no advantage from a RDHIS perspective, because the centralised data model at its core remains unchanged.

Similarly, H. Lee et al. (2020) have designed a blockchain implementation (using Ethereum) which “does not physically replace the electronic health record system, as most hospital information systems store detailed EMRs in a secure database on site” (H. Lee et al., 2020, p.3). Rather, it uses blockchain in a novel fashion by acting as an intermediary security and access control layer, to the underlying health record which still exists in a centralised location. Users are able to upload their full health record to this centralised network, and then use blockchain to permit sharing of the data to other stakeholders (whilst also making the data compliant with FHIR). This approach is promising, yet it still does not encompass the following issues:

2. Literature review

- Fully dynamic sharing of a health record, which is itself being added to constantly
- Fine-grained control over access to different types of data
- Rescinding shared access, or configuring time-limited access.

Finally, and most egregiously for a RDHIS, it appears to duplicate shared data in yet another centralised database.

Rajput, Li, Ahvanooey, and Masood (2019) improve on this somewhat by using the Hyperledger Fabric blockchain software to permit relatively dynamic access control of a users health record in the case of emergency access. Still, it appears to be predicated on a pre-existing centralised database of user data upon which to build a sharing framework.

A seminal study by Ekblaw, Azaria, Halamka, and Lippmann (2016) proposed a supposedly decentralised medical record management platform using blockchain. However, the blockchain component simply points to items in a centralised database to which the patient has no direct access of any kind. To be fair, this system was not *designed* with the aim of patient access or control in mind.

Following this, Dubovitskaya et al. (2017) demonstrated an approach to sharing medical imaging using Ethereum to point to records stored off-chain in a centralised database, and; Xia et al. (2017) proposed a blockchain approach to auditing and data provenance, again integrating with a pre-existing and centralised data store.

Geekyanage Don and Motalebi (2021) developed a prototype using Hyperledger and demonstrated the practical feasibility of basic data entry and retrieval functions using a blockchain, whilst Ding and Sato (2020) built a consortium blockchain using full homomorphic encryption to secure health data before persisting on the blockchain. This last approach offers the potential to actually store data *on* the chain, and potentially avoids the issues found in prior examples which simply point to existing centralised data stores.

de Aguiar et al. (2020) identify significant potential for blockchain in healthcare, but mainly in the domains of reliability and safety – for example, medication reconciliation, supply chain and quality monitoring. Nevertheless, they do identify the capability to put data ownership and control in the hands of the patient as a key opportunity with this technology, which has not yet been realised. More helpfully they identify some of the core technical challenges currently faced by blockchain in the healthcare arena:

2. Literature review

Table 2.6.: Blockchain challenges in healthcare. Adapted from de Aguiar et al. (2020).

Challenge	Description
Latency	Validating blocks is time-consuming. Bitcoin can only process 7 transactions per second, which also indicates a scalability issue (H. Wang et al., 2018).
Security	Can be compromised if an entity can secure a 51% majority in the network. Some authors have shown that nodes with less than 51% can still be dangerous either via ‘selfish mining’ (Eyal & Sirer, 2014) or ‘stubborn mining’ (Nayak, Kumar, Miller, & Shi, 2016) strategies.
Resource consumption	Where Proof of Work (PoW) is required, energy costs are very high.
Usability	Can be complex to build, manage and use.
Centralisation	Blockchain is decentralised in principle, but can end up being highly centralised depending on the chosen architecture.
Privacy	In public blockchains, there is concern that individuals can be re-identified from supposedly anonymous transactions (Wahlstrom et al., 2020).

Source code is freely available for a number of blockchain implementations, including Bitcoin (<https://github.com/bitcoin/bitcoin>), Ethereum (<https://github.com/ethereum/>) and Hyperledger Fabric (<https://github.com/hyperledger/fabric>).

2.3.3.2. Holochain

Holochain is a framework for building serverless distributed applications. Although development began prior to the rise of Blockchain, it has received far less attention and there is only a handful of peer-reviewed articles discussing it.

Holochain utilises an ‘agent-centric’ approach⁶¹, where all members of a distributed network can share data. Despite its name, it is architecturally much more akin to BitTorrent than Blockchain; it uses the same DHT approach to distribute data amongst users and perform consensus/validation. Since the DHT is central to Holochain’s operation, it is worth exploring how this works.

In Holochain, each user has their own ‘source chain’; this is a local database of all their personal data and actions they have taken in a network. By default the source chain is private

⁶¹Reflecting the recommendations of Bollinger et al. (2019).

2. Literature review

so, when that data goes offline (for example, if you shut down your laptop or switch off your phone), the data would not be available to anyone else⁶². Where the data needs to be shared (according to whatever rules you define), we therefore need a way to make sure that data is accessible beyond that single access point. This is achieved via the DHT⁶³.

The DHT is a public data store, where you have shared some metadata from your source chain (metadata is shared on the DHT by default) with a random selection of peers in the network. These peers validate and verify that data. Each piece of data has its own unique address, and the data is stored in the nodes (users) which are numerically closest to the address of that piece of data. The peers ‘gossip’ amongst themselves to find out who should store the data, and so that they know who to ask in future if the data ever needs to be retrieved. As an example, consider the scenario shown in figure 2.11. In this representation, we see a range of users connected in a Holochain network with each of them sharing some data which has been distributed across all users. The coloured lines demonstrate that no user has access to all the data but, collectively, the network is able to access all the public data. The overlapping of the coloured lines demonstrate that data is retrievable by any user *even where not all users are online*. Since the data is distributed across the DHT, it can persist and be available at any time (Holotescu & Vasiiu, 2020).

Holochain utilises the DHT by default, but the developer can configure exactly what information is publicly available in the manner shown in figure 2.11. It is easy to imagine that, in developing a healthcare application, it would not be desirable to make all health data entries public; this is a decision point that will be familiar to blockchain developers and is reflected in that literature. The data is either made private (known as a side-chain in blockchain), or a full entry is published to the public DHT in encrypted form. Whatever decision is made entails a fundamentally different workflow from a design and development perspective. For example, an app might be configured so that no data is publicly available on the DHT by default and therefore *only* stored in users’ source chains. However, an access control layer known as ‘capabilities’ provides the ability for users to pick and choose how they might share that data – with whom and for how long⁶⁴.

Janjua et al. (2020) utilised Holochain in a study proposing a proactive security forensics platform in the Internet of Things (IoT). Holochain was utilised by writing security logs to a DHT, which could also be verified and checked for non-repudiation. Some perceived advantages


⁶²The reader may argue that this seems like a form of centralisation, and this is technically true. However, the point of centralisation here rests with the individual that the data relates to and no other entity has decision-making authority over it.

⁶³Refer to Holochain documentation at https://developer.holochain.org/concepts/4_dht/ (accessed 7 May 2022).

⁶⁴Refer to Holochain documentation at https://developer.holochain.org/concepts/8_calls_capabilities/#remote-call (accessed 7 May 2022).

2. Literature review

Figure 2.11.: Holochain Distributed Hash Table



of using Holochain, according to this study, were that:

- A global leader consensus was not required, meaning that computing power was not wasted
- There is no dependency on typical consensus processes, such as PoW.

One of the key advantages, however, would be that of scalability. We have already seen that Bitcoin can only globally process up to 7 transactions per second and, furthermore, as the Bitcoin blockchain grows, the burden on each node in terms of storage and computation will only increase. Blockchain performance will degrade at scale, by design.

Conversely, Holochain's DHT approach means that performance only *increases* at scale – there are more peers with which to share data, which reduces the overall burden on the network (Frahata, Monowar, & Buhari, 2019). As Wahlstrom et al. note, “as a blockchain gets bigger, it gets less efficient and more resilient; but as a hApp [a holochain application] gets bigger, it gets more efficient and more resilient” (2020, p.5), in a manner similar to the efficiencies demonstrated by the BitTorrent protocol.

Furthermore, since holochain is serverless (by virtue of being wholly distributed), most hApps will run only on users' personal devices (for example, a laptop or phone) and therefore we can see that the concept really does support personal data sovereignty in a very explicit way. Finally Wahlstrom et al. (2020) note that holochain is much more consistent with the 'right to

2. Literature review

be forgotten’ privacy principle than blockchain where all data is immutable by design, and the full ledger is distributed across multiple users. To be clear, a Holochain app *is* immutable by default; the ‘chain’ part of Holochain is a sequentially ordered record of entries made by users. The difference is that hApp developers can configure this to work in a wide range of ways – from providing every user a full record of all public entries, to allowing users to ‘modify’ or ‘delete’ entries, giving effect to a ‘right to be forgotten’.

Holochain source code is published under a Cryptographic Autonomy Licence⁶⁵ at <https://github.com/holochain/holochain>.

2.3.3.3. m-ld

m-ld is a protocol for live information sharing amongst a distributed network. Its core principle is about keeping data as close as possible to where it is being used – rejecting traditional approaches which default to a centralised approach. In practice this means, for example, that a user’s health data would reside on their device but can be read or written to nearly instantly by any other actor in their network who has appropriate authorisation.

m-ld is currently at a very early stage in development but its approach and ethos is very much aligned with the principles of data sovereignty and distribution. Data being shared is referred to as a ‘domain’; in our case this might be an entry in a patient health record. This data can be replicated to other instances of the app (clones), and these are all then kept in sync by m-ld. It might be that a patient record is shared in this way with a patient’s care team, or their family. Anyone with access to that app instance can see new entries in real time without having to do anything else.

Data persistence in m-ld has parallels with Holochain when the DHT is not utilised. That is, users going offline can render data inaccessible unless a strategy is implemented to ensure that some persistent storage exists. This is all up to the developer in terms of how they architect and implement a specific app.

Since there is no access control functionality by default in m-ld, its role in a RDHIS would be either a set of bilateral relationships with a health professional (between patient and doctor, for example) or by using m-ld as a component in a broader application architecture. Authentication and authorisation are the responsibility of the app developer. However, the m-ld documentation also highlights an interesting feature which is of great relevance to this thesis. This is an audit function, whereby each clone maintains a journal of activity so that all actions can be audited⁶⁶. This is something that would need to be developed separately in Holochain and, of course, is a

⁶⁵A new form of licence, developed specifically for Holochain, and which specifically aims to ensure that the privacy of users’ cryptographic keys is protected.

⁶⁶Refer to m-ld documentation at <https://m-ld.org/doc/#security> (accessed 7 May 2022).

2. Literature review

key feature of blockchain.

Despite being relatively new, m-ld already offers a lot of promise for those who are interested in truly distributed architectures and it has the potential to deliver on key tenets of the data sovereignty movement – a well-designed and architected m-ld app would permit data sovereignty and, with some additional work, can facilitate the access control functionality which is crucial for a RDHIS.

m-ld source code is published under an MIT Licence at <https://github.com/m-ld>.

2.3.3.4. Solid

Solid attempts to solve the perceived over-centralisation problem found in the modern internet, where data and services have become concentrated in the hands of a few large corporations. It does this by separating data and identity. For example, on Facebook your data and activity is intimately tied to your identity (your account username which, in Facebook’s case, must be your real name⁶⁷). A user stores data in online ‘pods’ and can manage access to these pods, for applications, services or other users. There are two key concepts in Solid:

- **Pods.** Where your data is stored. A user may have data stored in multiple pods.
- **Identity provider.** This is a core part of Solid, which has recognised the need to manage identities and accounts. Identity providers can issue a WebID to a user which they use to sign in. This WebID is globally unique.

One important thing to note about Solid is that Pod providers and Identity providers will generally be third parties who provide that service to users⁶⁸. Some defaults that one would expect in such a system – such as the right to data portability – are not necessarily guaranteed in Solid, but are determined by the Pod providers’ terms of service. Furthermore, the Solid Project notes that the business models for Pod and Identity providers will be “determined by the market”, most likely with some charging a fee and others funded by advertising.

Parrillo and Tschudin rightly point out that, while Solid offers a *shift* in the power balance, it also “raises centralization risks insofar as barriers could be created to make the provider change [data portability] prohibitively expensive ... as well as a risk of an oligopoly dominated by the incumbents” (2021, p.1). It is not yet clear who Pod or Identity providers will be but, if it is indeed “left to the market”, then an optimal outcome for data sovereignty may not be

⁶⁷A policy which is marketed as making it easier for friends to connect with you, but has proven to be exclusionary where people with indigenous names have either had accounts suspended or not been permitted to use their real name in the correct format (A. Tan, 2015).

⁶⁸Refer to project website at <https://solidproject.org/faqs#the-business-model> (accessed 7 May 2022).

2. Literature review

forthcoming. I should also point out that pods can be *self-hosted*, although it is not clear what effort or complexity is involved in doing so at this stage.

A very appealing advantage of Solid is this separation of data from service and identity. A potential use case provided by Parrillo and Tschudin (2021) is an app which tracks a user's favourite movies. The data for that user (a list of their favourite movies) would be stored in a Solid pod, while the app would connect with other web services – for example, to obtain plot summaries and user ratings from IMDb. The user is free to access their own list and reuse it at any time; it is not locked in a proprietary service.

The practical argument against this use case eventuating is that apps would have to be designed with this model in mind. Given the acknowledged value of personal data, and the advertising revenue that can be harvested from it, it is very difficult to imagine Netflix, for example, taking the effort to fundamentally redesign their app so that user data can be stored separately. It is likely to be those very same market forces that dictate the success or failure of the Solid project.

2.3.3.5. Interplanetary File System (IPFS)

IPFS describes itself as a peer-to-peer protocol for storing and sharing data in a distributed file system⁶⁹.

Leveraging some of the technology we have already reviewed, IPFS relies on DHTs to break shared data into chunks and make it available to other nodes (users) who may want to access it. Whilst a Holochain DHT is unique to the particular app a user is using, however, IPFS aims to create a *global* namespace which means that any IPFS client can theoretically utilise all the advantages of a DHT to access information. At a global scale, this approach clearly has enormous advantages around redundancy and resilience and, furthermore, is independent of internet backbone connectivity (it can be used offline).

IPFS is already being used in a wide range of applications and has “hundreds of thousands” of nodes (Graber, 2021); Microsoft have used it as part of a decentralised identity project, Netflix use it internally for resource management, while support for the protocol is natively available in the Brave browser⁷⁰. Despite this, Parrillo and Tschudin consider that IPFS is still not user-friendly and “lacks compelling end-user applications” (2021, p.1).

An interesting development in this area is that IPFS also removes the ability for governments or corporations to restrict or block access to content. The DHT lets users host data (or parts of the data) themselves, and share access with any other users – there is no URL or single IP address that any entity could block to gain control of that content⁷¹. IPFS was used exactly in

⁶⁹Refer to IPFS website at <https://ipfs.io/> (accessed 7 May 2022).

⁷⁰Refer to Brave announcement at <https://brave.com/ipfs-support/> (accessed 7 May 2022).

⁷¹This also applies to Holochain, and any approach using a DHT.

2. Literature review

this way during 2017 when the Turkish government removed access to Wikipedia, proclaiming it a threat to national security. Activists quickly moved to copy the entirety of Turkish Wikipedia and host it themselves using IPFS, putting it out of government reach (IPFS, 2017).

The vision of IPFS, therefore, is more akin to a peer-to-peer internet. Because of its vision and mission, access control is not a core part of IPFS and can only be enabled by implementing encryption on shared data (Graber, 2021). This is not dissimilar from one of the approaches found in Holochain. However, Holochain's built-in 'capabilities' workflow is very aligned with personal health data management. IPFS's focus on distributing *public* data means that, while it is very exciting, it is unlikely to be the best fit for health data management without a large amount of additional work.

2.3.3.6. Scalable Secure Scuttlebutt (SSB)

SSB is a protocol, originating from Aotearoa New Zealand, that enables distributed peer-to-peer communication. Apps can be built using SSB which continue to work offline, or in high latency contexts, without any centralised components. While it was originally built as a basis for a decentralised social network, it has many potential applications.

Compared to a centralised model, where big tech host global pools of data that are delineated by the connections or choices of authenticated users, SSB starts from the insight that "each participant is only interested in a subset of the global data pool, thus it is feasible to locally store all the data a participant is interested in" (Tarr, Lavoie, Meyer, Tschudin, et al., 2019, p.1). Users therefore self-select the information that they care about (for example, by connecting with other users) and this defines the slices of data they store locally and can interact with.

Technically there is perhaps most similarity with Holochain, by virtue of an append-only single-writer log (functionally equivalent to a Holochain 'source chain') which contains a standard metadata set (for example, ID of the author, sequence number of the message, a creation timestamp, etc). Logs between connected peers are synchronised via a gossip protocol which performs replication. This, for example, ensures that new pieces of information are pushed to a users' log if they do not yet have access to them. In a similar way to both Holochain and blockchain, messages in a SSB network are immutable and not refutable (they are signed by the user).

Its conceptualisation as an alternative to social media means that "participation currently favors privileged users for whom privacy issues are not critical" (Tarr et al., 2019, p.9). What this means is that SSB is pseudonymous and, together with immutability and non-refutability, there is a possibility that data could be used to justify persecution. This risk can be addressed, however, where an SSB application is not 'public' and access control policies are well considered.

In summary, SSB has conceptual alignment with Holochain in many cases and considers itself

2. Literature review

wholly distributed; although it does not utilise a DHT, it fully supports offline distributed peer-to-peer communication. Its efficacy in a RDHIS application is limited somewhat by the absence of a proven access control method, although this is currently under development⁷².

2.3.3.7. Summary

Having reviewed a selection of the most viable candidate approaches for implementing a RDHIS, I will now summarise and compare their key features.

Blockchain has received a large amount of attention from researchers, and the literature base is enormous. This is likely to be due to its links with cryptocurrency, and the venture capital that surrounds that market. Decentralised or distributed apps that do not have such a clear link with monetisation have predictably received significantly less attention, making Twitter's 'decentralised ecosystem' review (Graber, 2021) a welcome change. The consequence of this, of course, is that valid criticism of the other approaches I have reviewed in this section has simply not had the opportunity to surface yet. This situation may therefore put blockchain at a relative disadvantage when comparing approaches.

At the very beginning of this research process, however, I should acknowledge that all of the content I had been casually reading led me to the belief that I would be building a prototype app using a blockchain technology. In this section we have reviewed the numerous criticisms and issues that surface in the literature – from problems with scalability (de Aguiar et al., 2020) and efficiency, to latency (Croman et al., 2016) and a disingenuous hijacking of the 'decentralised' narrative (Walch, 2019). On a personal level, I rapidly became disenchanted with how far blockchain really is distributed. It was this experience that led me to explore some of the other technologies found in this section.

Table 2.7 summarises all reviewed approaches against some key metrics that I have derived from the whole of this literature review and my initial understanding around what a RDHIS needs to do. This will be tested during research phase one (chapter 4).

An important issue we have not touched upon yet is that we must consider *how difficult is a technology to implement?* There are two dimensions to this. Firstly, I must consider complexity from the perspective of building a prototype for this thesis. This mainly encompasses what is within my technical acumen to complete, and is understood as a key consideration in the DSR literature (Gazem et al., 2018). Secondly, I must consider how any output of this research could be carried forward and scaled out into the real world. This will be discussed in chapter 8.

⁷²Refer to the ssb-tribes project at <https://github.com/ssbc/ssb-tribes> (accessed 7 May 2022).

2. Literature review

Table 2.7.: Comparison of candidate approaches

	Blockchain	Holochain	m-ld	Solid	IPFS	SSB
Latency	Potentially high	Low	Low	N/K	Low	Low
Scalability	Low	High	High	Medium	High	High
Implementation complexity	High	High	Potentially low	High	High	High
Ability to control access	High	High	Low	High	Low	Low
Serverless	No	Yes	Yes	No	Yes	Yes
Tendency to centralisation	Yes	No	No	Yes	No	No
Ability to audit transactions	Yes	Yes	Yes	Possible	N/A	Yes
Immutable data	Yes	No	Yes	No	No	Yes
Secure against malicious peers	Yes	Yes	Yes	Yes	Yes	Yes

What table 2.7 shows is some variability around specific measures, such as scalability and a tendency towards centralisation. Approaches that utilise a DHT are, by definition, considered to be fully distributed. Solid presents a categorisation issue by virtue of its aspirational messaging around decentralisation, but the technical detail shows that this is rather limited in practice (Parillo & Tschudin, 2021). At this early stage, non-negotiable requirements are that the selected approach must:

- Be highly scalable
- Provide access control
- Be genuinely distributed
- Provide an audit function.

2. Literature review

The other metrics are important, to be sure, but could be deprioritised in the context of a ‘proof of concept’ artifact such as I am proposing. Employing this list identifies Holochain as the most suitable technological approach at this stage.

3. Methodology and Methods

In this chapter I clarify the methodology used, and the research methods. The main aim of the thesis is to establish attitudes towards health data sovereignty, and to then demonstrate how it could be practically achieved. There are four core research phases:

1. I will establish the views of potential users around the ability to own and control their own health data
2. I will build a ‘proof of concept’ prototype demonstrating that it can be practically achieved
3. I will evaluate the prototype artifact
4. I will seek the additional views of relevant experts – to be chosen based on the findings from item 1, above – to locate the research in a broader policy context.

Whilst methodology is about the principles and processes that guide the practice of research, methods are the actual tools used to execute that research. The methodology must therefore align with the overall research aim. The chosen methodology is Design Science Research (DSR). The first part of this chapter explains the background of this methodology, and the reasons why I believe it to be appropriate. The second part discusses the methods to be utilised, which will encompass all four of the phases noted above.

3.1. Design Science Research (DSR)

DSR has been described as “a research activity that invents, or builds new, innovative artifacts for solving problems or achieving improvements” (Iivari & Venable, 2009, p. 3). It is predominantly about creating new things and expanding knowledge, and has been particularly popular in Information Systems (IS) and Computer Science (CS) research. It is very important to distinguish this from more routine design and build processes. A typical design process is to apply existing knowledge to existing problems, using best practice artifacts which are already part of the knowledge base. DSR, though, “addresses important unsolved problems in unique or innovative ways, or solves problems in more effective and efficient ways” (Hevner et al., 2004, p. 81).

3. Methodology and Methods

The same authors developed seven guidelines, outlining what may be considered a ‘threshold’ for DSR. These are as follows:

- DSR must produce an artifact
- There must be a technology-based solution to an important problem
- The artifact must be rigorously evaluated
- There must be a clear contribution to the research and knowledge base
- Rigorous methods must be applied to construction of the artifact itself, as well its evaluation
- The process determines an effective solution but, depending on the problem situation, may not empirically be the best
- The work must be presented effectively to a both technical and non-technical audiences (Hevner et al., 2004).

In the case of the research aims of this thesis, it would meet the DSR threshold in principle due to it being a largely unrecognised problem or, at least, one that has limited visibility because there is no apparent solution. The requirement to also contribute new artifacts to the knowledge base, of course, lends itself well to doctoral study.

3.1.1. Artifacts

DSR is supposed to be the bridge between IS research and practice. It is a methodology which introduces rigour to the practice component, as well as both grounding in, and contributing to, the research and knowledge base. It has both practical relevance (via the production of artifacts) and scientific rigour (via the formulation and utilisation of design theory). However, apart from one key component, there is still no broad consensus on what DSR should actually *include* (Peppers et al., 2006). That key component is that DSR must centre on *artifacts*, and this lies at the heart of DSR. In fact the perceived focus on artifacts is also seen to be a weakness by some researchers. Hevner et al. (2004) have expressed concern about the risk of overlooking the theoretical base by focusing on artifacts. Conversely, they also considered that this does balance the behavioural sciences, where there is arguably greater focus on theory perhaps at the expense of utility.

Artifacts can take one of four forms: constructs, models, methods or instantiations. This taxonomy has been in place since the publication of a seminal article by March and Smith (1995). Even in relatively recent literature (for example, Baskerville et al. 2018) it remains

3. Methodology and Methods

Table 3.1.: Artifact forms, as described in March and Smith (1995)

Artifact type	Description
Construct	Concepts and language where problems and solutions are defined
Model	Use of constructs to represent real world contexts of the problem and solution space
Method	Processes, such as algorithms
Instantiations	A tangible solution demonstrating feasibility

largely unchanged, reinforcing the point that artifacts are really a central tenet of DSR. An explanation of these different artifact types is shown in table 3.1.

Whereas March and Smith (1995) discuss these artifact types as being somewhat independent, later writers established a more holistic methodological connection between them. Hevner et al., for example, suggest that a successful ‘instantiation’ must “show that constructs, models or methods can be implemented in a working system” (2004, p. 79). That is, the first three artifact types shown in table 3.1 must be established in the research, but brought to life via the instantiation.

Baskerville et al. (2018) go a step further, and suggest that each artifact type must inform the next in a connected cycle. While the ‘construct’ will be derived from existing knowledge, to inform a better understanding of the problem and potential solution, the ‘model’ will rely on that construct to represent the problem and solution in a real world context. Finally, in what may be viewed as the final stage, the ‘instantiation’ brings all this together and demonstrates feasibility. This discussion should again reinforce that the artifact is the goal of DSR – whether that is inventing new artifacts where none previously existed, or improving existing artifacts.

3.1.2. Theory vs action

The tension between research and practice, noted earlier, is a key issue in the DSR literature. The basis seems to be that the very definition of an instantiation implies pushing the boundary of existing knowledge or theory; technological advances will necessarily *precede* theoretical or scientific advances. Designing useful artifacts is complex because of the need for advances in areas where existing theory is lacking; IS research is constantly advancing into new applications that were not previously encompassed in established theory (Hevner et al., 2004). Gregor and Hevner (2013) have described this tension by suggesting there are in fact two core groups within DSR:

- The design-theory group, emphasising design theory. For example, Gregor and Jones (2007), and

3. Methodology and Methods

- The pragmatic-design group, with an emphasis on artifacts. For example, Hevner et al. (2004).

As noted earlier, even this typology is difficult to make concrete since Baskerville et al. (2018) have reinforced the point that DSR has the potential to balance both groups quite adequately – and place an *equal* emphasis on theory and practice. They go on to note that, whilst few published DSR papers do effectively combine both groups explicitly, if an artifact is produced which truly does meet the entry threshold for DSR then it has by definition already contributed strongly to the knowledge base. This issue seems to come down to the interplay between science and technology.

Science is about descriptive knowledge of the natural world and behaviour, through application of a rigorous method. Technology, conversely, is about prescriptive knowledge of artifacts which have been designed to improve human capabilities – “new technologies are driven and enabled by science but, more often, scientific advances are driven and enabled by the emerging use of new technology” (Baskerville et al., 2018, p. 361). That is, technology (or, in DSR, the production of instantiations), does indeed make a significant contribution to the knowledge base almost by default. Technology can evolve very rapidly, whereas science takes a longer term approach to use of rigorous methods to understand *how* and *why*. Generally speaking, therefore, technology will *precede* science. But what does this mean for DSR?

The DSR practitioner must build a technological artifact, but provide generalisations and deeper understanding through nascent design theories expressed as models, methods, and constructs. Not only, therefore, does the instantiation derive from all other artifact types, but it generates more of them which science can then build upon to develop larger and more generalisable theory.

3.1.3. The DSR process

One of the most noticeable things about the DSR literature, is the amount of published articles specifying which steps should be taken and in which order. This is certainly never prescribed, but for researchers more used to a theoretical focus in the behavioural sciences, it can be disarming to read different versions of what appears to be a step-by-step process. This has probably arisen from the view of IS Researchers who complained that, while there is a lot of literature around how to do DSR, the “level of abstraction of those models is still too high in terms of providing a complete methodology” (Gazem et al., 2018, p. 2).

One of the original DSR articles, from 1995, asserted that there are really only ever two key stages to consider: build something, and then evaluate it (March & Smith, 1995). Certainly this is eminently practical and, in fact, sounds a lot more like something a developer, or a designer,

3. Methodology and Methods

Table 3.2.: Design Science Research Process. Adapted from Peffers et al. (2006).

1. Problem identification and motivation	Define the problem and justify the value of the solution. The problem definition should be used to develop the solution. Requires knowledge of the state of the problem, and the importance of the solution.
2. Objectives of a solution	Take solution objectives from problem definition. They can be quantitative or qualitative. Requires knowledge of problems and the efficacy of current solutions.
3. Design and development	Create the solution. Determine the desired functionality and architecture.
4. Demonstration	Show efficacy of the solution. This could be via experimentation, simulation or case study.
5. Evaluation	How well does the solution solve the problem? Compare the solution objectives to observed results from the demonstration. Can go back to step 3 if required.
6. Communication	Communicate the problem, the problem importance, and the built artifact to relevant audiences.

would do every day in the course of their work. In other words, a process that is perhaps not very *methodological*. One key difference, of course, is that DSR should not deal with routine design problems; the threshold is greater in that it must deal with an unsolved problem in an innovative way. But the process described by March and Smith (1995) does certainly seem to fall into the ‘design-theory’ group discussed in the preceding section.

Peffers et al. (2006) describe a DSR process which focuses on dissemination and profiling of a DSR project, which has six steps. These are shown in table 3.2. This approach has a focus on communication of the research, something else that has been identified as problematic in the DSR field. This, again, seems to sit firmly in the ‘design-theory’ group, via an implicit deprioritisation of theory. It is also clear from table 3.2 that it possibly is too abstract to follow as a blueprint for successful DSR research. A helpful balance is struck by Gazem et al. (2018) who describe a very comprehensive fourteen point process which appears to effectively blend theory and action. This model is described in table 3.3.

This DSR Roadmap clearly has more steps, and more detailed ones, than seen previously in table 3.2. From the perspective of a researcher, the first four are particularly useful. It is good practice, especially in doctoral study, to capture and be clear about where the thesis idea came from, and to be systematic about determining relevance and feasibility. Furthermore, the explicit steps around defining requirements and making preparations are useful practical tips. The balance with design theory comes from the genesis of the idea (in steps 1 and 2), the dissemination of the completed work (step 14) but, most importantly, from checking the knowledge base prior to development (step 10). At this stage, the researcher would have a broad

3. Methodology and Methods

Table 3.3.: Design Science Research Roadmap. Adapted from Gazem et al. (2018).

1. Document the spark of an idea/problem	Explain where the idea came from. This might be a researcher seeing a new way to solve an existing problem based on their own knowledge, or from other academic work. Some DS researchers contend that a problem should generally come from the real world rather than literature (for example, Benbasa and Zmud 1999).
2. Investigate/evaluate the importance of the idea	Researcher needs to make sure that the problem is: unsolved, important, and that it contributes to a knowledge base. Hevner et al. (2004) recommends empirical research to do this step.
3. Evaluate the new solution feasibility	The research must ensure the solution can actually be achieved. This is not simply a proof of concept issue, but there must also be adequate budgets and time available. Subsequently, does the researcher have the technical knowledge or competence to understand how to build solution, or to implement it.
4. Define research scope	Define the objectives and the research scope. How do we limit the research to what we know is feasible?
5. Resolve if study is within Design Science paradigm	The research must meet core DSR principles, for example of requiring the development of artifacts, or of iterating between building and evaluating. Furthermore the artifact must not be so specific to needs of a single entity that it cannot be generalisable (even if generalisation and extrapolation does not occur within the original DSR project itself).
6. Establish type (IS Design Science vs IS Design Research)	Winter 2008 found that DSR has two main types: IS Design Research, which is the work of constructing and evaluating artifacts, and; IS Design Science reflects Design Research and creates standards for its rigour, more like the bigger research process around the doing. The current research fits into the Design Research paradigm because it seeks to create a solution by using a build and evaluation process. This step feels a lot more like the two groups discussed by Gregor and Hevner (2013) in section 3.1.2.
7. Resolve theme	Is the research about Constructing, Evaluating or Both? (Hevner et al., 2004) has lamented the number of studies that work on construction with little or no evaluation.
8. Define requirements	Determine the necessary resources, tools and skills to complete the research. First define the main tasks that need to be achieved, then identify the fundamental requirements to meet them – what are the requirements, what are the instruments/methods, what are the outputs?
9. Define alternative solutions	It is important to assess a wide range of possible solutions. This can be a product of the researchers own creativity, and domain knowledge. This point is also referenced in (Hevner et al., 2004).
10. Explore knowledge base support of alternatives	This links the existing knowledge base with what it is proposed to build, so the approach can be justified.
11. Prepare for design/evaluation	Prepare a plan for building the solution discussed in step nine. This should also include defining criteria for performance, as well as determining how evaluation will be carried out.
12. Develop	This Roadmap focuses on setting up the project on a sound basis, and does not prescribe artifact construction methods which should be used. This is expected in DSR since it relies on the “creativity of the researcher to choose the methods or techniques to construct the artifacts” (Gazem et al., 2018, p. 8).
13. Evaluate	What elevates work from routine practice to academic is the extent and rigour of evaluation. This can take two forms: Artificial evaluation – where feedback is gained from experts, or; Naturalistic evaluation – where the artifact is implemented in a real world case study. It is possible to do both types at a single evaluation stage, before going back to step 12.
14. Communicate findings	This is about dissemination and publication of findings, as well as artifacts (particularly constructs and models).

3. Methodology and Methods

grasp of what is possible, and what the theoretical basis is of prior relevant work (if it exists). Finally, the theoretical component in step 14 – where artifacts other than instantiations can be disseminated, and built upon – should not be underestimated. When executed well, this is indeed the bridge between research and practice that DSR set out to provide.

3.2. Research approach

Bearing in mind the above, we can start to think about a broad approach for this research. I will use the DSR Roadmap, as described by Gazem et al. (2018), as a framework.

3.2.1. Problem provenance, rationale and significance

Firstly, the problem identification/relevance is discussed in section §1.1. The same section also touches on the importance of the research idea, whilst chapter 2 confirms the idea's status as an important and unsolved problem. Something that I have mentioned previously is that I personally feel the issue is important, and I can verify it is unsolved. For the general public, it is likely it is not viewed as an issue but this may be because there is currently no technological solution to it and, hence, it goes entirely unaddressed. The scale of engagement from a sample of the general public will form a core research phase in this thesis.

3.2.2. Feasibility

Evaluating feasibility (step 3) is a core part of DSR and the building of instantiation artifacts – there is a pushing of the known technological boundaries, to contribute something new to the knowledge base (Baskerville et al., 2018). However, Gazem et al. (2018) expand this to encompass other feasibility issues, such as time and technical competence. Time and budget constraints can be addressed further by refining the scope of the study (step 4). Prior to designing a prototype, I have no reason to suspect there will be any feasibility issues. If there is a concern it would be my lack of domain experience with distributed technologies in general, and the overcoming of any learning curve required to utilise a suitable candidate such as Holochain. I note this as a concern primarily to reflect the point I had previously made around the centralisation hegemony. Technical approaches and mental models have embedded a centralised approach, and it is not simply a programming language or a syntax quirk that has to be learnt; it is a whole new understanding around data and information.

3.2.3. Scope

Regarding step 4, the initial research scope has already been constrained. Most importantly, I have elected to remove any development work around end-user products. That is, anything to do with a service, app or website that a non-technical end user might access. The logical end game for this research is that a RDHIS is developed in such a way that anyone can access it and use it. Doing this requires significant work in app development, user interface design, accessibility and translation services – amongst many other factors. This is not only outside of my core technical competence, but would be untenable to achieve in the context of a thesis. The scope is therefore reduced to simply establishing that core functional requirements can be met (these are discussed in section §5.2).

3.2.4. The DSR threshold

Is this study definitely eligible for a DSR approach? Referencing back to the seven principles set out by Hevner et al. (2004), the answer is yes. When carried out as planned, there will be a rigorously evaluated artifact (or artifacts) that solves an important problem. It will be making a clear contribution to the research and knowledge base and, by virtue of this thesis (and potentially other publications), the findings will have been disseminated to a range of audiences.

3.2.5. Define requirements

Definition of requirements is a critical step (Gazem et al., 2018). The DSR researcher must first identify the main tasks required to complete the study, and then break these down into instruments and requirements. In my case, the main tasks required to achieve what I aim to do in this research are:

1. Understand attitudes to personal health data sovereignty in Aotearoa New Zealand.
2. Develop knowledge of Holochain.
3. Develop prototype instantiations.
4. Assess applicability and performance of prototypes against user attitudes and requirements (evaluation).
5. Understand perspectives of domain-specific subject matter experts (SMEs) to develop an implementation roadmap for a RDHIS.

Based on these tasks, I can break these down further and identify the requirements, as shown in the table below.

3. Methodology and Methods

Table 3.4.: Study requirements. Based on Gazem et al. (2018).

Requirements	Instruments	Outputs
Data collection	Survey / interviews	Understanding of user attitudes and requirements, and subject matter expert (SME) perspectives
Design evaluation method	Experimental/simulation	Final prototype meeting requirements
Experimentation/learning	Technology/software	Developed prototype
Data analysis	Software	Analysing collected data and using findings

The next layer of requirements is utilised in the prototype design and evaluation stages, and will be derived both from the literature review and the user survey.

3.2.6. Define possible/alternative solutions

This step is essentially a horizon scan for all possible solutions, to ensure that the most generally effective one is being progressed in the DSR project. There is an overlap here with Guideline 6 from Hevner et al. (2004); ‘Design as a search process’. This step exhorts the researcher to do their best to assess possible solutions, whilst recognising that it will not be realistic to find the best possible solution in most valuable IS problem spaces – “the search for the best, or optimal, design is often intractable for realistic information systems problems” (Hevner et al., 2004, p. 88). This step is therefore at the problem-solving heart of DSR – how do we effectively solve whatever unique or tricky problem is before us?

Simon (1996) has written about problem solving in the IS space, and proposes a balance of *means*, *ends* and *laws*. Means are the tools, resources and actions available to build a solution for the problem; ends are the DSR objectives, and; laws are uncontrollable factors operating in the solution space. Again, Hevner et al. (2004) notes that one need not be absolutely comprehensive in identification of these three factors; DSR can legitimately abstract some of these away, or focus on a simple set of *subproblems*. It is most critical to have a useful starting point upon which subsequent research can build. For example, even within a single DSR project, the specific ‘laws’ or ‘means’ at play may not be apparent to the researcher until one or several prototypes have been constructed. A second, or even third, prototype will therefore have uncovered a lot more knowledge about these variables, and this already makes a valuable contribution to the knowledge base.

In terms of scope – if the researcher is assessing all known solutions, how far should they go? It is not practical to enumerate all solutions, and assess *everything*. Simon (1996) introduces the concept of ‘satisficing’, which involves a focus on heuristic search approaches. In practice, this

3. Methodology and Methods

means that it is appropriate for the artifact to simply “work well” for that problem space – “the critical nature of design in IS makes it important to first establish that it *does* work ... even if we cannot completely explain *why* it works” (Hevner et al., 2004, p. 90). For this thesis, the focus is on building a solution that does work against specific known outcomes. The aim, in brief, will be to provision a health record which is held by the individual, can be securely shared and audited, and is without any centralised components. It will simply either work, or it will not work. Therefore, there is little chance of the procedural ambiguity noted in Hevner’s discussion on heuristic search strategies.

What seems most logical is to leverage technology which is already explicitly about distribution, and has that principle at its core. We have already selected and reviewed a representative cross-section of technologies in section 2.3.3 and summarised their relative merits. As part of that discussion, we established that both Blockchain and Solid are decentralised models, whereas the remainder can be considered fully distributed (Baran, 1964). The summary to that section concluded that Holochain appeared to be the most effective candidate technology and therefore it will be utilised to build the prototype.

3.2.7. Explore knowledge base support of alternatives

In this step, the researcher must link the existing knowledge base with the proposed instantiation artifact. This will be done initially in chapter 2, but subsequent chapters will make this link more explicit as the detail of the proposed solution is discussed.

The literature review encompasses a broad range of issues related to the data sovereignty problem. We established that data sovereignty can have multiple meanings and is heavily mediated by social context. For example, it is a focus for indigenous groups due to its potential to mitigate the impacts of colonisation. We then reviewed some international examples and saw that no jurisdiction is currently investigating distributed solutions; a global trend, in fact, is towards geopolitical localisation. We also reviewed the current state of distributed/decentralised technologies and linked it with the data sovereignty concept to theorise that it can provide a solution. The DSR process which is described in chapter 5 will attempt to formally link the problem with a possible solution by way of instantiating an artifact.

3.2.8. Prepare for design/evaluation

In this stage we are required to formally plan how we will carry out the design and evaluation of the artifact. This, as well as the subsequent two ‘Develop’ and ‘Evaluate’ steps, is discussed fully in chapter 5 and chapter 6.

At this stage, no design can be formalised until the first phase of research has been completed.

3. Methodology and Methods

Similarly it is difficult to specify appropriate evaluation methods. Nevertheless, the scope of the study implies that some evaluation methods will not be suitable. For example, I have already clarified that I am not focusing on a product which is ready for end user utilisation as an output of this research. This automatically removes a range of possibilities around user acceptance testing, or user-centred design. Being primarily code-based, the evaluation is likely to be a simulation of some kind – to demonstrate simply whether or not the requirements have been met. For example, it will be possible to define some key functional requirements and then to automate tests which can be formalised as part of the final codebase. This is discussed more fully in chapter 6.

3.3. Research phases

I have already noted that there are four streams of research activity planned in this project. These align with the foregoing discussion about DSR in general, and the research approach drawn from the DSR Roadmap presented by Gazem et al. (2018). These four phases, together with methods and outputs, are shown in figure 3.1.

Since each of these phases follows in a logical sequence, it is important that each phase is discussed individually in more detail. The following chapters will step through each research phase, before culminating in a Discussion (chapter 8).

3.4. Ethics Approval

All research undertaken for this study was subject to the AUT Ethics Guidelines¹. Specific ethics information related to different research phases is summarised below.

3.4.1. Consumer Panel Survey

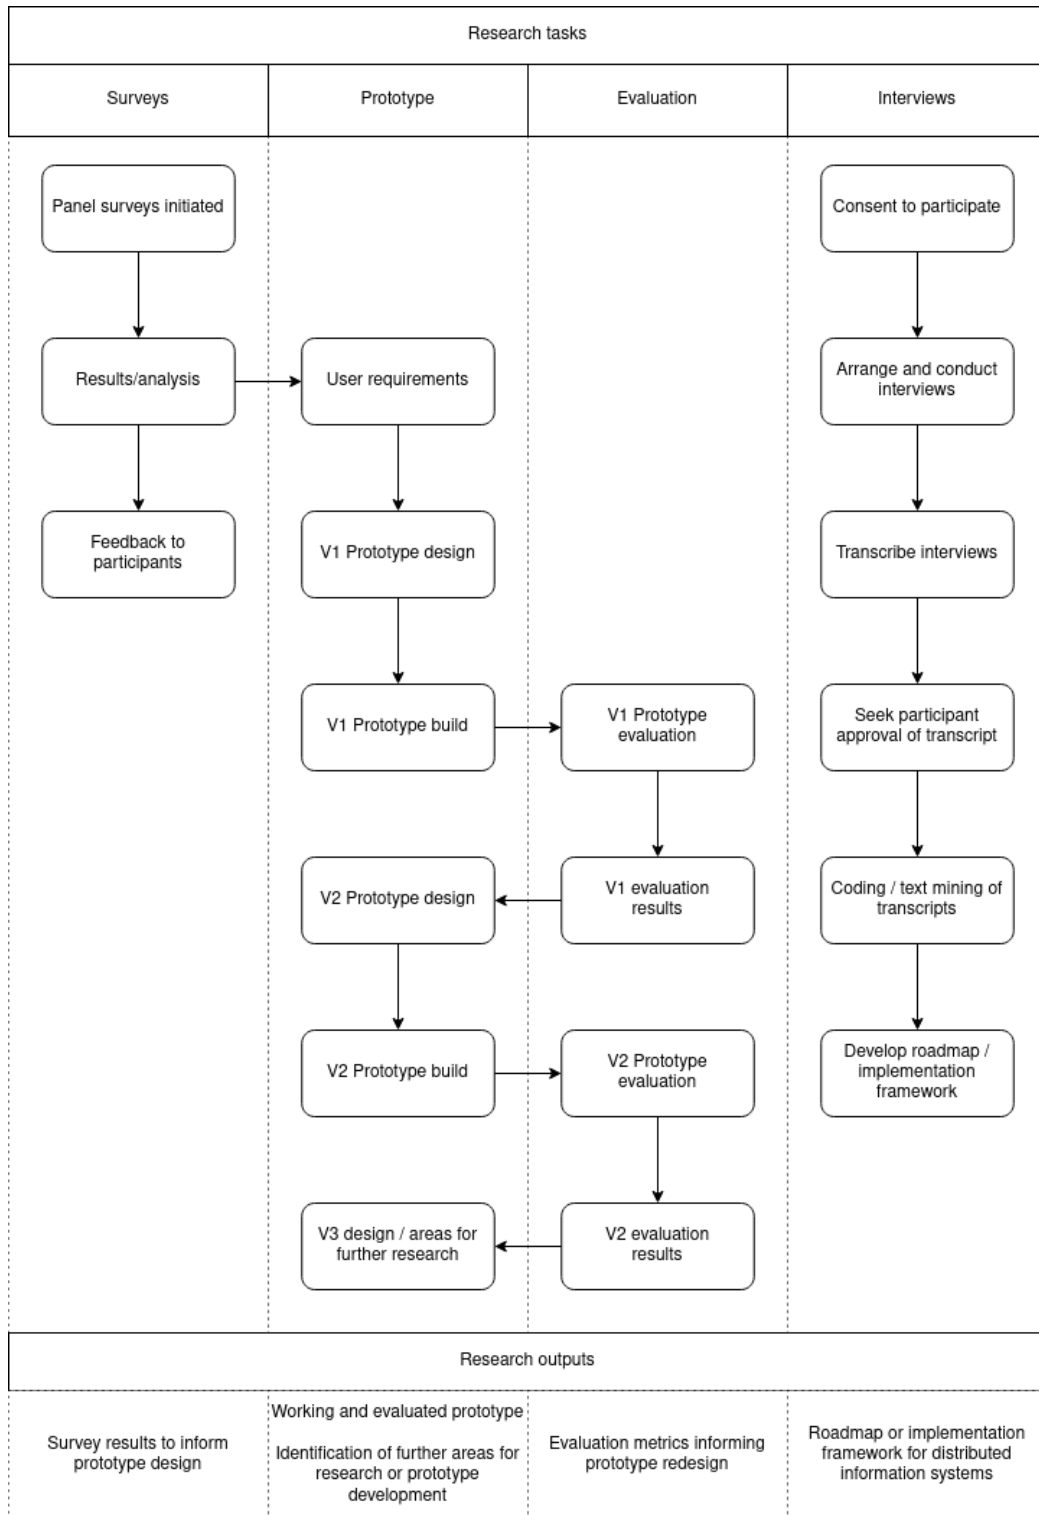
The first phase of primary research involved a survey of the Midlands Health Network (MHN) Consumer Panel. This is discussed in more detail in chapter 4.

The research protocol submitted to the AUT Ethics Committee, and given full approval on 10 July 2020, specified that all participants would be required to give informed consent to take part, and how the information would be utilised. The Participant Information Sheet for the survey was directly accessible on the landing page, and is reproduced in full as appendix A. Participants are also required to formally indicate consent prior to accessing the survey itself – this is captured in the response. The approval letter is available in appendix B.

¹ Available at <https://www.aut.ac.nz/research/researchethics/guidelines-and-procedures>.

3. Methodology and Methods

Figure 3.1.: Research phases and outputs



3. Methodology and Methods

As a follow up to the survey, summary information was released back to the respondents via MHN and also provided to governance bodies within the organisation for their reference.

3.4.2. Interviews

The fourth phase of research involved a series of semi-structured interviews, and this is discussed in chapter 7.

The research protocol was given full approval on 17 May 2021, and clarified that all participants are required to give informed consent to take part. All participants were provided with the Participant Information Sheet (or the opportunity to review and ask questions) and this is reproduced as appendix E. The approval letter is available in appendix D.

4. Research Phase 1: Consumer Panel Survey

In order to establish some of the core DSR principles – as set out in Hevner et al. (2004) and Gazem et al. (2018) – it is necessary to begin with conducting primary research into user attitudes and requirements. Some extant research around this topic has already been discussed in section 2.2.3. The goal of this research phase, from a DSR perspective, is to establish the following:

- Is it really a novel problem?
- How valuable or important would it be to solve it?
- Do users have specific requirements that can inform prototype design?

This will be achieved via an online survey of a pre-existing health consumer panel, which is discussed in detail in section §4.6.

Why was the survey method selected? For discussion around survey methodology and methods for the survey component, I will rely on the authoritative ‘Survey Methodology’ (Groves et al., 2009). Whilst surveys are perhaps one of the oldest and most well known research approaches, they are especially suited to certain problems or objectives. For example, Groves et al. note that they are especially capable at measuring attitudes and opinions, and particularly useful for those interested in “what people knew, felt, and thought” (2009, p.4). This is exactly the goal of this research phase, where I am trying to understand opinions around ownership and control of health data – as well as hypothetical scenarios intended to elicit how they would utilise a RDHIS if it did exist.

In terms of survey design and execution, Lavrakas (2008) notes that we must start with some important foundational elements, to ensure we have clarity and understand exactly what information we are seeking from respondents. In fact there are nine key stages to effective survey design and execution and these are shown in table 4.1. I will deal with these nine stages in the following sections.

4. Research Phase 1: Consumer Panel Survey

Table 4.1.: Stages of survey design. Adapted from Lavrakas (2008).

Stage	Description
1	Determination of goals and objectives
2	Definition of key concepts
3	Generation of hypotheses
4	Choice of survey mode
5	Question construction
6	Sampling
7	Administration and data collection
8	Summarisation and analysis
9	Conclusions and communication of results

4.1. Goals and objectives

Having clearly defined goals helps to centre and focus the survey, and acts as a critical framework for construction of questions. Each component of the survey should contribute directly to realisation of the identified goals. Whilst I have already identified thesis research questions, in section §1.3, this is different since here we need to identify the objectives specific to this *survey*. The survey is merely the first phase of a DSR process where first we have to ascertain if this is indeed an unsolved problem that can be solved in an innovative fashion.

I would like here to refer back to the typology presented by Gazem et al. (2018), and visualised as table 3.3. Using that framework we are really only dealing here with ‘Step 2’, which is about investigating and evaluating the importance of an idea. That is, the survey needs to verify that owning and controlling one’s own health data is actually a problem that people feel needs to be solved. Furthermore, I have identified that the survey would be a valuable opportunity to also ask some questions around exactly how a RDHIS might work in practice. This in turn can inform the later ‘Step 8’ – defining requirements for the prototype. I can therefore summarise the survey goals as follows:

- Measure the potential demand for ownership and control of health data
- Measure the potential demand for greater insight into how health data is utilised
- Understand more about exactly *how* people might want to share their health data.

4.2. Definition of key concepts

I have identified and discussed some key concepts in section §2.1. I was conscious that in the actual design of the survey itself, some care would need to be taken around how these terms are utilised. After testing some discussion of concepts and terminology with friends and family, I determined that it was best to omit any jargon where possible. I therefore did not use terms like ‘data sovereignty’, or even ‘privacy’, in the survey itself. These terms were replaced by plain language description. For example, my initial draft survey asked a question about attitudes to “ownership and control of health data”, rather than the concept of data sovereignty or the technical components around distribution.

4.3. Generation of hypotheses

Having conducted an initial literature review, I was reasonably confident I had an understanding of people’s attitudes and opinions around access to data and data sharing. As noted in chapter 2 there is a pattern in the literature that, whilst people can be fatalistic about data capture, there is generally a sense that they want more control when they are asked about this directly (Aitken et al., 2016; Garrison et al., 2016; Moon, 2017). I carry this hypothesis through to this research phase, and reasoned that people may not be agitating for greater ownership of health data for a couple of reasons. Firstly, there is the fatalism around not being able to control their personal data which has been noted by McMullan (2015) and Solon (2018). Secondly, this is compounded by a lack of understanding about technical solutions which offer greater ownership and control of personal data.

In this way the survey has to speak to these hypotheses, and we should be able to conclude that: people do actually want ownership and control of health data, and; they can imagine how to use that power if a technical solution was available to them.

4.4. Choice of survey mode

For some research objectives the selection of survey mode is obvious. For example, when researching literacy or access to the internet you should not use an online tool. In most other cases then, the researcher is required to balance pros and cons in making a decision (Groves et al., 2009). Due to the chosen sampling method, and what is essentially a pre-existing constraint on the operations of the MHN Consumer Panel, I actually had very few options around survey delivery mode. Most MHN surveys are delivered via the SurveyMonkey tool. The Panel members are used to this survey mode, and it is the one they would expect. In keeping with the principles

4. Research Phase 1: Consumer Panel Survey

of DSR it was most efficient to simply leverage this established process, in order to access some results and insights to inform prototype design. Nevertheless, it is worth briefly discussing some of the perceived merits and issues associated with self-administered web surveys so that, as a researcher, I can acknowledge error or bias when using the results.

In general terms, online surveys have the baseline coverage error issue that only people with access to the internet can use them. Furthermore, there are accessibility issues to note for people with physical disability or vision impairment. The two key advantages are:

- Measurement error can be reduced, since research has shown that people are more likely to answer truthfully when they are not being interviewed in person
- Online surveys tend to be the least expensive survey mode (Cowles & Nelson, 2015).

I should also add that, while the easiest choice is to continue using an online survey mode since it is what the sample population will expect, there is then another layer of decisions to be made about exactly how that is done. We will look at design and construction of questions next, but also I have specified some technical components around the survey design and delivery in section §4.7.

4.5. Question construction

This area is thought to be particularly critical to reducing measurement error in surveys. Groves et al. note that “questions that are easily understood and that produce few other cognitive problems for the respondents introduce less measurement error than questions that are hard to understand or that are difficult to answer for some other reason” (2009, p.259). Intuitively, this feels rather obvious however it can be difficult to implement in practice, particularly where the researcher is a subject matter expert in the research topic and some of the key concepts may not be a common part of the lexicon of the sample or target population.

In order to minimise survey measurement error, three standards must be met as shown in table 4.2. Evaluation of a proposed survey must ensure that these standards have been addressed, and there are a range of mechanisms to achieve this – each with relative advantages and disadvantages. Reflecting the constraints around what is practical in a thesis, as well as the DSR focus on artifact development, I have opted to use what Lavrakas (2008) refers to as ‘conventional pretesting’. This is where a small number of surveys are deployed, allowing the researcher to quickly identify problems around delivery and interpretation of questions. I shared the draft survey with four of my friends and family, as well as the MHN Consumer Panel co-ordinator, and asked for their feedback around the three standards. This is summarised in table 4.3.

4. Research Phase 1: Consumer Panel Survey

Table 4.2.: Survey question standards. Adapted from Groves et al. (2009).

Standard	Description
Content	Do the questions ask about the right things?
Cognitive	Do respondents understand the questions, and are they able to answer them?
Usability	Can respondents complete the survey easily?

Table 4.3.: Conventional pretesting results

Standard	Issue identified in pretesting
Cognitive	Confusion around use of jargon such as ‘data sovereignty’
Cognitive	Need to clarify exactly what is the scope of ‘health data’
Cognitive	Some of the response options provided are ambiguous
Usability	Text is overlapping on certain screen sizes
Usability	Response options are not easy to see on mobile devices

It should also be noted that the group of people conducting this basic pretesting were not diverse enough to draw any conclusions around how, for example, the cognitive standard would be assessed by different ethnicities or people with different first languages.

The results however pointed to some general issues around layout of the survey itself, as well as a need to clarify wording and terminology. As already noted, this ‘cognitive’ standard is critical to a well constructed survey; it introduces the concept of *reliability* – which is the extent to which you can expect different groups of people to understand and answer the question in the same way. For example “which city do you live in?” is a highly reliable question, whereas “how much money does it take to be happy?” is not (Cowles & Nelson, 2015). I elected to increase reliability via two methods: structure of the questions as statements, and; removing all jargon or language that was not plain. In this manner, for example, I moved away from a question asking “what are your views around health data ownership and control?” to a statement – “I want to be able to own my health data, and decide who can access it” – where respondents are asked to weight their agreement via a Likert scale.

Another important consideration in question construction is *validity*. That is, are you measuring what you think you’re measuring? Having clarified survey objectives (described in section 4.1) I felt confident I knew exactly what I was trying to ascertain. The selected questions were centred on those objectives. Some confidence that I had done this reasonably well was provided by a lack of ‘content’ standard issues reported back from pretesting.

4. Research Phase 1: Consumer Panel Survey

Cowles and Nelson (2015) identify three key components of good question construction. These are summarised in the following table:

Table 4.4.: Key components of effective question construction. Adapted from Cowles and Nelson (2015).

Description	
Specificity	Does the question have <i>validity</i> ?
	Congruence between objectives and the question
	Relevant to the respondent
Clarity	Does the question have <i>reliability</i> ?
	Gap between researcher and respondent knowledge is minimised
	Minimal jargon or technical terms
Brevity	Questions are straightforward and uncomplicated
	Complex sentence structures are avoided

Some other areas to be cautious of are: loaded questions, double-barrel questions, and double-negative questions. I made sure that these were absent where possible, or minimised. Let us therefore look at the survey structure as it stood after the pretesting phase:

1. Tell us about yourself:
 - Select your age group
 - Select your gender
 - Select your ethnicity group
2. How should a health information system work? Indicate how far you agree with the following statements (options are: Strongly disagree; Disagree; Neutral; Agree, and; Strongly agree):
 - I want to be able to own my health data, and decide who can access it
 - I want to be able to record and manage my own health data (eg. from smart devices or wearables) and see it alongside all my other medical information
 - I want to be able to see who is using my health data
 - There should be a way for health workers to access my data in case of an emergency
3. Access to my health data. Imagine you could control all of your health data. Indicate how you would allow access to it, for each of the following purposes (options are: On an ongoing basis; Only as needed or upon request; Only in an emergency, and; Never):
 - I would give access to my GP..

4. Research Phase 1: Consumer Panel Survey

- I would give access to a hospital team...
- I would give my partner, family/whānau or caregiver access...
- I would give the Ministry of Health information about my hospital visits...
- I would give other government agencies access (eg. MSD or ACC)...
- I would allow my data to be used in a government database, so it can be combined with other government data and used for policy analysis and research...

4. Any other thoughts? Please write down any other thoughts you have here.

As can be seen, most of the survey features structured or closed questions. While this does introduce the risk that the researcher biases the responses via design of the structured options (Krosnick & Fabrigar, 1997), it also does allow much greater uniformity and the ability to clearly gauge consensus on issues. Structured questions must be *exhaustive* and *mutually exclusive* (Cowles & Nelson, 2015).

The Likert scale referred to in question 2 encompasses every possible response type that could be offered, and none of them overlap. The second response range, shown above in question 3, is more complex since the options are somewhat more abstract. The respondent is being asked to comment on the basis on which they would share data, across time. While the options are mutually exclusive, it is probable that they are not exhaustive. To compensate for this, I offer the four options as the ones I consider to be most obvious, following a literature review and my own knowledge of the subject, however the respondent can also add extra free text comments against that question if they wish ¹.

4.6. Sampling

One of the key issues around surveying is to understand how generalisable the results are, and this is where the concept of sampling takes a prominent role. The survey methodology literature places emphasis on representation and generalisability, depending on the formal objectives of the research project. Nevertheless, the researcher will usually want to be able to extrapolate from their sample to draw broader conclusions. The ‘target population’ for this research is ostensibly any adult who has accessed a health service in Aotearoa New Zealand². But, of course, it is not feasible for me to survey this target population and so I must use a sample.

Whilst Fink notes that “a good sample is a miniature version of the population of which it is a part ... The best sample is representative, or a model, of the population” (2003, p.2), they also

¹40.6% of respondents offered any kind of unstructured response; 6.6% of all respondents left an unstructured response against this question specifically.

²Note that while the age bands referenced in the results in the following sections include 15-24 years, there are no respondents or Consumer Panel members under 18. This age group is used simply to retain consistency with established age groups used by the Ministry of Health and Statistics New Zealand.

4. Research Phase 1: Consumer Panel Survey

caution that “no sample is perfect” (2003, p.3) and that there will inevitably be some level of bias. Broadly speaking there are two types of sampling – Probability and Nonprobability. I have utilised Nonprobability sampling, and this is thought to be appropriate in three main scenarios:

1. Surveying hard to find groups – for example, people engaged in types of criminal activity
2. Surveying specific groups – for example, people in hospital due to specific diseases or conditions
3. Surveying for pilot situations – for example, assessing performance of a prototype or pilot initiative (Fink, 2003).

Clearly the third item above fits very naturally with the DSR approach, where the goal is to be agile and move relatively quickly to the design and build of a workable artifact. On this basis nonprobability sampling was considered to be appropriate for this research phase. Specifically I have utilised the ‘Convenience sampling’ method, which is simply a “readily available group of individuals or units to be used” (Fink, 2003, p.17). Whilst this approach may seem inferior compared with probability sampling, and its ability to offer more rigorously scientific conclusions, it essentially becomes a cost-benefit decision driven by the research objectives.

Designing and implementing a probability sample-based survey would not have been impossible for this research phase, but would have added significant time and cost overhead which could not easily be justified within the constraints of a thesis – more so given that the methodological focus is production and evaluation of an artifact, rather than a statistically pure survey. The obvious downside to this approach is that representativeness of the sample is unlikely to align with the target population, and therefore the capacity to generalise may be limited. Nevertheless Lavrakas contends that, as long as this limitation is recognised, convenience sampling has an important role in certain types of research and does “allow some quick exploration of a hypothesis that the researcher may eventually plan to test” (2008, p.150).

The researcher’s personal network led to identification of a pre-existing Consumer Panel, based in an Aotearoa New Zealand Primary Health Organisation (PHO). I received approval upon application to conduct my research with the Midlands Health Network (MHN) Consumer Panel from the organisation’s Digital Governance Group and the CEO. The panel is composed of patients enrolled with that PHO, and who are surveyed on a monthly basis around topical issues to do with healthcare and the business of the PHO. This Panel has been in place since 2018 and, since that time, surveys have been conducted on: Health information and privacy; Patient portal utilisation; Pharmacogenomics, and; access to healthcare services. This Panel is formally set up as a group of people who can be called on for their views, to steer strategy and planning. At the time of writing, the Panel numbers 1,670 individuals in total and MHN indicate that the response

4. Research Phase 1: Consumer Panel Survey

Table 4.5.: Consumer Panel gender

Gender	n	%
Female	1163	69.6
Male	504	30.2
Other	3	0.2

Table 4.6.: Consumer Panel age group

Age group	n	%
15-24	45	2.7
25-44	295	17.7
45-64	744	44.5
65+	586	35.1

rate is generally around 30% on average. Therefore, I might expect a total of ~500 responses to a survey, if it is conducted well and captures the interest of Panel members.

All Panel surveys are completely voluntary and the ‘convenience sampling’ method was widely used by the PHO to generate insights about patient attitudes and opinions. I was cautioned that the Consumer Panel was not necessarily representative of the Midlands population, and it is worth specifying more detail about this group for context. When we break the panel members down by key demographic variables we can see that older European females dominate, as shown in table 4.5, table 4.6 and table 4.7.

These tables appear to be unrepresentative of Aotearoa New Zealand as a whole, when we consider the published Census 2018 ethnicity rates of 71.8% European and 16.5% for Māori

Table 4.7.: Consumer Panel ethnicity

Ethnicity	n	%
Asian	34	2
Māori	131	7.9
MELAA	5	0.3
European	1482	88.7
Pasifika	18	1.1

4. Research Phase 1: Consumer Panel Survey

(Statistics New Zealand, 2020)³. However we should also consider some regional variation amongst those national figures.

The equivalent locality figures for the Waikato District Health Board (DHB) region (nominally the main coverage area of Midlands Health Network – although it also encompasses Taranaki and Tairāwhiti) shows both a higher proportion of Māori (22.8%), and a lower proportion of Pacific peoples (3.1%) when compared to Aotearoa New Zealand overall (Ministry of Health, 2019). On this basis it is fair to say that both Māori and Pasifika are particularly under-represented in this panel, and the survey results should be viewed in this context.

4.7. Administration and data collection

As already noted, the MHN Consumer Panel is used to conduct online surveys and that is the selected survey mode. The actual administration of the survey was subject to some constraint, since I was essentially piggybacking off the organisation's established process. That process is to issue one online survey at the beginning of each month. The survey remains open for the rest of the month, after which a new one is issued. Results and summary data are sent to the Panel usually around two weeks after the survey has closed⁴. My survey was made available to the Panel on 4 September 2020; an email was sent by MHN to the Panel members providing a link to the survey, and some brief introductory text that I supplied. The survey was open until 30 September 2020 at which point it was closed.

Survey topics are anything that is topical or of interest to the organisation at that time. There are occasionally staff members using this avenue for survey deployment if they are completing independent research for study. Naturally, all topics must be broadly relevant to the strategic direction of the PHO. My application was accepted on the basis that it is to a large extent about patient empowerment, and also parallel issues that are very topical to do with privacy and data sovereignty.

On the topic of privacy, I felt it important that the research role models, as far as possible, the commitment to information privacy. With that in mind, the survey is completely anonymous. There is no requirement to capture any identifying information, although technically the collected data still meets the threshold to be considered as 'personal information' in the Privacy Act 2020. Whilst name or email address is not part of the survey, it is necessary to collect some basic demographic variables (these are gender, ethnicity group, and age group – as shown in section §4.5). The purpose of collecting this data is twofold: firstly to understand how repre-

³Please note that direct comparisons with Census ethnicity data are difficult, because it is possible to select multiple items as part of the Census and hence totals do not add to 100%.

⁴Therefore, a maximum of six weeks for feedback in the event that a respondent completed it on the day the survey was opened.

4. Research Phase 1: Consumer Panel Survey

sentative the Consumer Panel respondents are, and; to identify any key variances which can be attributed to one or a combination of those demographic variables. Whilst we already know that the baseline Consumer Panel is not generally representative, it will be instructive to measure how the profile of the respondents compares.

Another issue to note around administration is that none of the questions were mandatory, with the exception of the initial consent. This enables people to only answer the questions they wish to answer. However, there is a counterargument that some people may answer none or very few questions (leading to incomplete data), or, some people may inadvertently click through the whole survey and complete it without having supplied any data. I consider these last two risks to be acceptable, when compared with enforcing the answer of some questions and potentially upsetting or disengaging respondents.

The major downside of an anonymous survey is that, if a user wished to withdraw their information or correct it (as mandated by the Privacy Act 2020), this would be nearly impossible. As a mitigation for this, the survey tool itself makes it very clear that any submitted data cannot be altered or removed – since it will be impossible to tie any response to an individual. Making this point very explicit at the beginning of the survey process is the most effective counter-balance to this issue.

On a more technical note, and in keeping with the principles behind this research, I have designed and built the survey tool itself entirely by hand to give confidence that there is none of the tracking or fingerprinting that is becoming more ubiquitous in use of the internet via the rise of surveillance capitalism (Cyphers & Gebhart, 2019; Zuboff, 2019). In technical terms the tool is a JavaScript application (using a third-party SDK by SurveyJS⁵), built as a static website in the Amazon Web Services (AWS) S3 object storage service.

The site incorporates the Information Sheet completed for the Ethics Committee application, as well as direct contact links for myself and the supervision team. The survey tool was accessible at a URL and then taken offline once the survey closed. Survey responses were stored in the JSON format and were made available for querying and analysis only to myself within the AWS platform.


4.8. Summarisation and analysis

There are two components to analysis of the survey data. Firstly, quantitative analysis has been conducted on the structured question responses. Secondly, qualitative data analysis was conducted on the free text unstructured questions (of which there were three). These will be discussed separately.

⁵A free open source library available at <https://surveyjs.io/>

4. Research Phase 1: Consumer Panel Survey

Figure 4.1.: Ethnicity comparison of Consumer Panel and survey respondents



4.8.1. Quantitative analysis of structured question responses

The MHN Consumer Panel has 1,670 members in total, and all were invited to take part in my survey. 426 members completed the survey, giving a response rate of 26%. Since we know that the demographic breakdown of the Consumer Panel is not representative (as reviewed in section §4.6), the respondents to my survey are also not going to be generally representative.

Firstly figure 4.1 shows that the ethnicity mix of the survey respondents was notably different to the Consumer Panel overall, with a slightly higher proportion of European respondents and lower proportions of Māori, Pasifika and Asian respondents.

Regarding gender, figure 4.2 shows that the survey received a higher proportion of male respondents and a lower proportion of female respondents. My survey also received a total of 1.1% Decline responses or simply not answered for gender.

Perhaps the biggest difference between my survey and the Consumer Panel is the age profile, shown in figure 4.3. My survey had a significantly lower proportion of respondents in the 25-44 and 45-64 age groups, and a corresponding higher proportion of respondents in the 65+ age group.

4. Research Phase 1: Consumer Panel Survey

Figure 4.2.: Gender comparison of Consumer Panel and survey respondents





Figure 4.3.: Age group comparison of Consumer Panel and survey respondents



If we take the survey respondents and bring all these demographic variables together, we can find the five largest subgroups. This is shown in table 4.8.

The first set of survey questions (formally in section two, after the demographic questions) asks the respondents to indicate their level of agreement with a series of statements. The results of this are summarised in table 4.9.

What the results above appear to show, very broadly, is support for the concepts of health data

4. Research Phase 1: Consumer Panel Survey

Table 4.8.: Top five respondent subgroups

Subgroup	% of total
European females aged 65+	28.4
European males aged 65+	21.4
European females aged 45-64	20.9
European males aged 45-64	8
European females aged 25-44	5.9

Table 4.9.: Responses to survey – section two

	Strongly dis- agree %	Disagree %	Neutral %	Agree %	Strongly Agree %	EMPTY %	Mean	SD
I want to be able to own my health data, and decide who can access it	9.9	2.6	13.6	38.5	32.6	2.8	3.8	1.2
I want to be able to record and manage my own health data	9.4	5.4	20.4	39.7	21.4	3.8	3.6	1.2
I want to be able to see who is using my health data	9.4	1.4	7.8	34.7	43.4	3.3	4.0	1.2
There should be a way for health workers to access my data in case of an emergency	8.9	1.9	2.1	38.5	45.3	3.3	4.1	1.2

4. Research Phase 1: Consumer Panel Survey

ownership and control – as well as auditability (being able to see who is using the data). There was less support for the idea of being able to record and add your own health data into the owned record. Interestingly, there were meaningful proportions of people responding neutrally to the first two questions. Whilst both had good support (71% and 78% respectively either agreed or strongly agreed), it is perhaps an issue related to a cognitive standard (Groves et al., 2009) that was not fully met. Or, they may simply feel neutrally about the topic. Probing this area can be done by conducting qualitative analysis of the free text responses in the following section.

Expanding on the quantitative component, the five response options were assigned a numeric value and these were used to calculate the mean and standard deviation. In the above table we can see that all items, on average, are supported by the respondents. However each item had a standard deviation of 1.2 which, on a five point scale, indicates a meaningful distribution of opinion. This is made obvious when we note that each question had a nine to ten percent response of ‘Strongly disagree’. In summary, there is broad support but a minority who seem opposed to change.

When these responses are broken down by demographic variables, there are some interesting points to note. The item ‘I want to be able to see who is using my health data’ provoked the most diversity; by age, the 65+ group agreed⁶ with this item at a much higher rate than the younger combined 15-44 year age group (82% versus 24.2% respectively). We can guess that this may be a product of older age groups having more experience of using the health system, and more opportunity to feel frustrated around access to health information. Belfrage et al. (2022), however, found that self-reported health status was one of the biggest influences on health system trust⁷ and unfortunately that data is not available from this survey. There was also a slight difference by gender, with 80.2% of females agreeing with this item compared with 75.2% of males.

The second set of questions, posed in section three of the survey, asks respondents to imagine that they really did have ownership and control of their data. In this scenario, when and how would they choose to share the data with different groups? The results of this section are summarised in table 4.10.

Section three is really aimed at informing requirements for a prototype design. Given that section two demonstrated support for ownership and control of data, I feel validated that I took this line of questioning in section three since I now have data available on how people would use a RDHIS. Since this part of the survey utilises only nominal data, I present only basic descriptive statistics above.

⁶This is a simple combination of ‘Agree’ and ‘Strongly Agree’.

⁷A potential explanation put forward by the authors is that people with better health have fewer contacts with health services and there is therefore less opportunity for trust to be broken.

4. Research Phase 1: Consumer Panel Survey

Table 4.10.: Responses to survey – section three

	Ongoing %	On request %	Emergency %	Never %	BLANK %	Mode
I would give access to my GP ...	89	9.2	0.5	0.5	0.9	Ongoing
I would give access to a hospital team ...	59.4	35	4	0.5	1.2	Ongoing
I would give my partner, family/whānau or caregiver access ...	48.6	34.3	12.4	1.9	2.8	Ongoing
I would give access to the Ministry of Health ...	30.3	48.4	11.5	4.7	5.2	On request
I would give other government agencies access ...	12.4	43.7	21.1	15.5	7.3	On request
I would share my data anonymously with a government database, so it can be used for policy analysis and research ...	38.3	34.3	8.5	14.6	4.5	Ongoing

4. Research Phase 1: Consumer Panel Survey

Firstly the obvious point to make is that there appears to be very high levels of trust in the respondent's GP, reflected in 89% saying they would have ongoing access to personal health data. The other items see greater diversity in response, and this again is validation of the concept that fine-grained access control of health data is really necessary since people have very different opinions about who they should share data with, and on what basis.

Following the high level of trust in GPs, both hospital teams and family/whānau receive good levels of support for ongoing access to health data. The notion of providing data to government appears to split opinion the most, with 30% saying they would offer full access for hospital teams and 15.5% saying they would *never* share data with any other government agency. By demographic variable, the item around family/whānau access received some of the most mixed responses. Willingness to share data with family/whānau on an ongoing basis is positively correlated with age; only 30% of 25-44 year olds would give ongoing access, compared with 54.8% in the 65+ group. The 25-44 age group also had the highest rates of answering 'Never' to this question (6.7%).

This item also has a key difference when viewed by gender, with males far more supportive of ongoing access compared with females (64.1% versus 39.9% respectively). This pattern is reinforced by age group, with males aged 65+ the highest supporters of ongoing family/whānau access at 71.7%. This is likely to be related to social dynamics around health, with research in the US finding that 61% of unpaid caregivers are women (AARP, 2020). Males also continued to support ongoing access to health data for other items at higher rates than females.

Finally, the last item provides an interesting spread of responses. This was actually intended to gauge opinion around Statistics New Zealand's IDI but, in the spirit of the cognitive standard outlined by Groves et al. (2009), I was required to provide a brief explanation rather than use its actual name (which most New Zealander's would not recognise). As can be seen from the data, it received a relatively high proportion of respondents selecting 'Never'. However, it also received a higher proportion of 'Ongoing' than the Ministry of Health.

It is possible that the keywords utilised against this item (that it is anonymous, and that the data would be used specifically for analysis and research) introduced a response bias whereas no such context was given to how data might be used by the Ministry of Health or other agencies.

However, it should be remembered that I am not necessarily trying to demonstrate a statistically accurate picture of trust in different parts of government. I wanted to ascertain whether a RDHIS requires dynamic and fine-grained access control, and indeed it does. This is a key finding that will inform requirements for the prototype.

4.8.2. Qualitative analysis of unstructured question responses

Having reviewed the quantifiable data from the survey, we must also look at the value held in the ‘free text’ responses available. In total 173 respondents offered a total of 227 free text responses and so, because this qualitative data is not particularly deep or rich, I have elected to use a content analysis approach⁸. Content analysis follows many other qualitative approaches in relying on the organisation of content into categories or keywords (Mackieson, Shlonsky, & Connolly, 2019).

Coding is a technical task, with the aim of deriving meaning from unstructured data. For example, all 227 responses can be tagged, categorised and then analysed in aggregate to gain new insights. Saldana notes that coding “enables you to organize and group similarly coded data into categories ... because they share some characteristic” (2015, p.8). Since there were a relatively small number of unstructured responses, I chose to take a very simple and practical approach.

Firstly, I read through all the comments in detail several times. This gave me a sense of what terms might make good codes or categories. Just to be clear, the focus here is retained on the objectives set out in section 4.1; although some comments were interesting, they were not necessarily relevant to the objectives.

The next step is to run through the comments and assign coding categories to them. This of course is rather subjective, but Saldana (2015) notes that this coding process should ideally be done over several cycles to repeatedly test and evaluate. For interviews or large quantities of qualitative data this process is completed using specialised software (for example, the proprietary NVivo). However, since I am dealing with a very small amount of data⁹ in this case I elected to simply code this by hand using very basic and ubiquitous spreadsheet technology.

To provide an example, the comment “*I thought my GP had access to my Indici Health Data at all times*” was coded with the single ‘Interoperability’ tag. This comment represented a reasonable proportion of responses who, rather than commenting on ownership and control, noted that they really just wanted their data to be easily available across the entire health system (and, in some cases, had specific examples of how surprised they were when they discovered that was not currently the case). I coded this as ‘Interoperability’ because this is arguably a parallel thread of issues around the broader health system being able to join system effectively across sectors, organisations and IT systems.

A second example is the comment “*I think my data is mine, and my DR and team can use as required, and if I land in the emergency room. Otherwise I should be asked for access.*”

⁸Reflexive thematic analysis is utilised in chapter 7 because that process generated a large amount of rich qualitative data.

⁹A grand total of 7,960 words, forming an average of 35 words per response.

4. Research Phase 1: Consumer Panel Survey

Table 4.11.: Top ten key concepts from First Cycle coding

Tag	Frequency
Control	64
Interoperability	45
Access	32
Ownership	20
Empowerment	19
Correction	11
Privacy	11
Healthcare	9
Security	9

(If the Dr has to send it to a specialist etc that would be okay too)“. I coded this with three tags – ‘Interoperability’, ‘Control’ and ‘Ownership’. The comment clearly calls out a desire for ownership and control. Finally the last statement in parantheses indicates again that there should be more joined-up data sharing across the health system, which speaks to the ‘Interoperability’ concept.

Following this first cycle, a total of 244 tags were applied across the 227 responses; 16 responses were allocated three or more tags; 53 were allocated two; 86 were allocated one, and; the remainder were not allocated any tags¹⁰. A master list of these tags was compiled, and summarised by counting the frequency of each. The top ten are shown in table 4.11.

This list gives us a quick insight into what the major themes of the unstructured comments were, and this helps to build a First Cycle coding process (Saldana, 2015). Conceiving of this process as a cycle is also noted by Coffey and Atkinson who suggest that coding “is usually a mixture of data summation and complication ... breaking the data apart in analytically relevant ways in order to lead toward further questions about the data” (1996, p.29).

On review of this First Cycle, one thing that did become apparent was that people were talking about the same concepts but from very different perspectives. Using the second comment example again from above, it is clear that the respondent is not simply mentioning the concepts of ownership and control but is *in favour of them*. Crucially there are several comments which also mention ownership, but indicate they are not supportive of it. I would need some way to separate these out, so I revised the list of tags to include ‘sentiment’. That is, where a positive or

¹⁰These were technically valid comments but did not offer any contribution to the core concepts identified. One example being “for me this is very important”.

4. Research Phase 1: Consumer Panel Survey

Table 4.12.: Final coding concepts used

Code	Description
Access	The respondent refers to a generic wish for greater access to their data
Agency	Sentiment about sharing data with government agencies
Audit	A desire to track who has been using their health data
Control	Sentiment about controlling health data
Correct	A desire to be able to correct information in their record
Effort	A notion that accessibility should be prioritised and made easy for non-technical users
Empowerment	A belief that data ownership, access or control can empower patients and lead to better health outcomes
Healthcare	General sentiment about the healthcare system
Interoperability	Refers to a focus on simply enabling the health system to share data wherever it is needed
Ownership	Sentiment about ownership of health data
Privacy	Respondents specifically referencing privacy issues or concerns
Security	Respondents specifically referencing security issues or concerns

negative opinion is expressed against a concept. In the final analysis, I retained twelve coding concepts for use against the unstructured data and these are summarised in table 4.12.

Another component of qualitative coding is that of *attribute coding*, otherwise known as *descriptive coding* (Richards, 2009). This is essentially about adding other variables into the coding process so that comments and concepts can be analysed by them. A good example in an interview process, is to add attribute codes for demographic variables, or perhaps other items such as their iwi/hapū or number of children. This is particularly important when interviewing, since the canon the researcher is working from is the transcript and it is generally not enough to simply say, for example, that 30% have text coded against a particular item. The researcher would feasibly want to know additional context, such as whether or not those respondents were employed. Saldana refers to this as “good qualitative data management” (2015, p.70).

4. Research Phase 1: Consumer Panel Survey

While this is a separate activity when interviewing, fortunately the survey did capture three key attributes – the demographic items of gender, age group and ethnicity. It is possible that more attributes could have been added, and played an important role in analysis, however I have endeavoured to work by the recommendation in the literature that parsimony is extremely important (Cowles & Nelson, 2015). In terms of analysis, then, let us review what we have found from the coding process.

- *Conflation of key concepts.* The fact that 45 respondents commented on ‘interoperability’, and 32 commented on ‘access’, suggests either that people were misunderstanding the true meaning and potential of ownership and control, or were simply using the opportunity to talk about other issues important to them. These two concepts represented a good proportion of the responses and I found, looking through the comments, that some respondents appeared to default to a position where this discussion was simply about allowing them to access and read their own health data. This perhaps relates back to the cognitive standard outlined by Groves et al. (2009), however the language used in the survey was unambiguous. It is also possible that some respondents found the concept too fanciful to fully countenance, and simply interpreted those terms according to their reality which will be influenced by the centralisation hegemony. This notion is alluded to in one of the comments: *“The concept of data ownership is rather abstract – I’d want to be sure that suitable people create data relating to me, within a structure that allows me to understand and influence access”*.
- *Patient empowerment.* This was referred to by 19 of the respondents – a minority, certainly, but some of the comments were extremely positive and hopeful about what this could offer in terms of empowerment. One example being: *“I really like the idea of accessing and controlling my own data and who can see it ... Sharing with whānau would be revolutionary for how chronic conditions are managed”*.
- *Sharing with government agencies.* This has already been identified as a topic with diversity of opinion in section 4.8.1. While 13 respondents mentioned the issue, the sentiment coding further revealed that nine respondents felt negatively about this, and four felt positively. Examples of each are: *“Government have a bad record at keeping information secure. They should never be allowed to access my data without my express permission”*, and *“Generally happy about sharing my health data to appropriate agencies on a need to know relevance [sic] basis”*.
- *The ability to see who is using my data.* This item was the subject of a specific question in the survey. In section 4.8.1 we saw that 78.1% of respondents agreed that being able

4. Research Phase 1: Consumer Panel Survey

to audit data usage was important. Only eight people offered unstructured comments on this topic, however. Two comments representing different sentiment on this are: *“I must know who, when and why my data is accessed”*, and *“I don’t think people need to see who is accessing data as laws already restrict inappropriate access to health information”*.

- *The ability to correct my data.* Whilst this notion is actually a core component of Privacy legislation in Aotearoa New Zealand (including the Health Information Privacy Code) it is interesting that this was a key issue for some respondents, perhaps indicating that this legislative right may not be working well in practice. Some key comments on this concept are: *“I would love the ability to actually correct the stuff that is written by the hospital. I find the amount of errors in the hospital notes I have seen utterly appalling [sic]. Things like the wrong year for surgery and wrong medication”*, or *“I would love to be able to correct information that was bulk loaded by my GP practice ... [which] made it look like I’d had a miscarriage in my 2nd marriage 12 years after it actually happened”*. The ability to edit information was not originally in scope for design of a RDHIS, but it seems clear that there must be a better feedback loop to review information entered in a persons health record.
- *General state of the healthcare system.* The idea that unstructured comments can become a dumping ground for pet issues was realised in this domain. Several respondents took the opportunity to pass comment on the broader health system in Aotearoa New Zealand. Of the nine respondents commenting on this concept, only two were positive. Examples of both sentiment are: *“Each visit I get the feeling – hurry up you are taking too much of my time. Here take a pill this may fix it, if it doesn’t, make another appointment. For me, a visit to the Medical center has become something I wish to avoid and I normally come away frustrated”*, and *“I feel that I am well served by the NZ Health Department”*.
- *Security.* This of course is a key part of any technical solution, although I have not discussed security in any depth yet. None of the survey questions either referenced this concept specifically, so it is interesting that nine respondents still spoke about it, one summing it up thus: *“I would not like my data to fall into the wrong hands or used against me”*. This is a concept that would have been formally mentioned in prototype requirements and design, however it is unlikely any comprehensive work can be done on this huge area in the constraints of this thesis.

4.9. Conclusions and communication of results

Summarising the above discussion, we can come to the following broad conclusions:

4. Research Phase 1: Consumer Panel Survey

- People really want to exert control over their health data, and see who has accessed it
- Respondents are not so interested in integrating health data from different sources (e.g. wearable devices)
- People are generally happy to share all their data with their General Practice
- Respondents want more control over sharing data with the wider health system, and their family
- There is moderate support for sharing data with government in a deidentified fashion
- The lowest level of support was for sharing identifiable data with other government agencies (specifically MSD or ACC).

Please note that these conclusions are based on a respondent group who are overwhelmingly comprised of older age groups, and European ethnicity. A full set of results, incorporating the above conclusions, was shared directly with the respondent group via the organising team at MHN. A copy of this report is available at appendix C.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

Referring back to figure 3.1, we can see that there are two key tasks for this phase of the research:

1. To convert the findings from Phase 1 into a formal set of requirements, outlining what the prototype must be capable of doing. There is a process of refining results, and interpretation, during this part and it will be discussed fully.
2. To build a prototype such that it can be evaluated against the requirements.

5.1. Defining requirements

In chapter 3 I discussed DSR in some detail. This methodology has a focus on building novel *artifacts*, but ideally in a way that is not simply a developer's flight of fancy. The bulwark against this is the attention that should be paid to clarifying objectives and requirements. The framework presented by Gazem et al. (2018), in fact, has *eleven* discrete steps before any prototype work is even carried out and clarifying requirements is specified as a particularly important one.

Different DSR projects will be generating requirements from a range of sources. We have already reviewed (in chapter 4) the survey carried out as Research Phase 1; the focus of which was to establish several of these precursor steps. Specifically this Phase 1 addressed the following steps from Gazem's framework:

- Establish that the problem is unsolved, important and it contributes to a knowledge base. Importantly the survey approach accords with the exhortations of Hevner et al. (2004) that this should be achieved by empirical research.
- Define the requirements, to identify exactly how the prototype should be designed and what tools or resources will be required.

To briefly take a step back from our specific requirements in this project, it is important to understand what kind of a prototype we are contemplating. Schork and Kirchner lament the fact that “the prototyping process itself is mostly driven by intuition and strongly depends on the

5. Research Phases 2-3: V1 Prototype design, build and evaluation

knowledge of the developer ... leading to a rather inefficient and tinkering-based process” (2018, p.2) with varying levels of overall success. This preparatory work is therefore very important.

Conversely, to take a step forward, we must understand that requirements are a critical part of the evaluation phase (discussed in chapter 6). Evaluation can only be carried out against a clear and concise set of requirements, that are ideally gathered via empirical research. In fact we can go further and assert that “a design artifact is complete and effective when it satisfies the requirements and constraints of the problem it was meant to solve” (Hevner et al., 2004, p.85). Requirements are indeed very important.

However, it is important to firstly recognise that, whilst the survey addressed the two steps specified above (problem definition and defining requirements), it was predominantly focused on the former. That is, the survey’s main objective was to *identify that ownership and control of health data was indeed an important problem that is unsolved*. Some specific design requirements certainly fell out of that process. However, ‘requirements gathering’ is a design and IS discipline in its own right and is the subject of a large body of research and literature. I am not undertaking a full-blooded requirements gathering exercise, and the DSR methodology does not require that I do. It is enough that I do not simply depend on intuition and my own tacit knowledge (Schork & Kirchner, 2018). So, what do the survey results tell us about requirements that we can use as an evaluation framework for the prototype?

5.1.1. Controlling access to data

Survey respondents indicated that they did want to exert control over their health data and wanted to decide who could access it. We can build on the findings from the survey, and require that access control is dynamic – that is, the user can permit and revoke access at any time, to any party, and this takes immediate effect¹. Therefore, a key requirement could be formulated thus: *Users must be able to freely permit or revoke access to their data at any time.*

5.1.2. Auditing data access

There was strong support in the survey for the ability to easily see who has accessed the shared health data (a mean score of 4 out of 5, and standard deviation of 1.2). While this was seen as an important factor in a RDHIS, respondent comments implied that people would be very happy to access this data within the current centralised ecosystem. Given the relative immaturity of distributed data as a concept, it is difficult to quantify how far this perspective is a general lack of understanding – perhaps itself resting on the conceptual dominance of centralisation.

¹Note that the revocation of access may not necessarily mean that the grantee no longer has access to any shared data. Two relevant scenarios are: where the grantee has copied data or stored it locally, or; where the data has been distributed across many peers. This distribution of data is discussed further in section 5.3.4.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

Furthermore, if there were genuine apathy – even in a scenario where decentralisation is well understood – then this may be a product of the demographic profile of the survey respondents².

This requirement can be formulated as: *Users must be able to easily see who has accessed their data and when.*

5.1.3. Correcting data

An unexpected theme from the survey was that respondents really welcomed the chance to be able to correct their health data in cases where it was erroneous. This is a tricky area in a RDHIS; the draft concept is that the vast majority of data will be made available to an individual, directly from source systems. This is technically simple to achieve. But the ability to directly correct data implies a different workflow, and potentially entails the possibility to write back into a source system. This seems unlikely (but not impossible) on a number of levels³.

We should also be aware of concerns that stakeholders already have with distribution, around the potential for users to falsify information upon which others may rely. Distribution certainly does not imply that data entries can be modified at will⁴ and, in fact, many distributed frameworks rely heavily on the principle of *immutability*. Where a record needs to be edited or updated, the original data is left in place but consumers can be pointed to a new or updated version of that record. This preserves auditability and permits inspection of changes that have been made to data which, in fact, aligns well with the audit requirement we have just discussed above.

More realistic is the consequence of simply having more people viewing their own data; genuine errors are noticed more quickly, and these could be flagged as part of the current provisions under the Privacy Act 2020⁵. This in fact provides a convenient parallel with our discussion around Open Source Software earlier in section 2.3.2.1, where Raymond asserted that “given enough eyeballs, all bugs are shallow” (2001, p.30).

For simplicity, at this stage I will assume that there is no requirement for users to be able to edit *recorded* data – but the access and auditability requirements will act as an important foundation on which to base correction requests under legislation such as the Privacy Act 2020 and the Health Information Privacy Code 2020.

²Specifically, we might expect that marginalised groups would care more about the overarching power and control dynamics that govern their personal data and are a unique selling point of real distribution. Unfortunately these were poorly represented in the survey.

³There are both technical and political obstacles to writing into incumbent source systems, and also legal issues around record keeping and data lineage.

⁴Please see section 7.7.3.2 for a discussion on data integrity, as part of the interview research phase.

⁵Principle 7 specifically provides the right for an individual to ask an organisation to correct their personal information, if they believe it is wrong. The organisation is not obliged to comply.

5.1.4. Security

Finally, a strong theme from responses was that users wanted any system to be very secure⁶ – not only from external intrusion, but to be safe from any potential misconfiguration whereby data was inappropriately shared. So this encompasses two facets of security – the app will perform as intended (its internal configuration does not permit unwarranted data sharing), and the app is hardened against external compromise (it implements ubiquitous security standards and approaches, such as SSL for data in transit).

However, the very core functionality of a RDHIS is that access to data can be controlled dynamically by the user. This is not just a desirable feature, but it also amounts to a key security consideration in the two following ways:

- *“I want to be certain that no one can access data that I have not specifically granted them access to”*
- *“If I revoke your permission to access my data, I want to be certain that you can no longer access it”.*

These are both security tests that should be run and verified as part of the prototype evaluation. The component about hardening from external compromise is certainly important, but will not be a priority for this proof of concept. It is something to consider in real world applications and is noted as an area for further research in section §8.3.

I will therefore define two new requirements:

- *No data can be accessed without appropriate permission*
- *No access to data is possible once that permission has been revoked.*

5.2. Prototype requirements

Having distilled the key requirements, we can now start to specify the things that an effective prototype needs to do. These will incorporate all the requirements that were specified in the previous section, as well as encompassing other core functionality that wasn’t specifically drawn out from the survey process.

These requirements will be quite specifically defined, and will be utilised as formal tests that will be run against the prototype, to gauge efficacy. As Hevner et al. have said – “a design artifact is complete and effective when it satisfies the requirements and constraints of the problem it was

⁶This was particularly interesting given that no survey questions directly referenced this topic.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

Table 5.1.: V1 Prototype requirements

Category	Requirement ID	Requirement
Recording and retrieval	RP1	A range of health data can be posted to Alice’s record.
	RP2	Only Alice has read access to this data by default ⁸ .
	RP3	Alice can review the data being posted to her record.
Access control	RP4	Alice can share her data with any other stakeholder, at any time ⁹ .
	RP5	Third parties with delegated access can review the shared data.
	RP6	Alice can reliably edit or revoke the delegated access at any time ¹⁰ .
Audit	RP7	Alice can review audit data, showing delegated permissions and details about third party use of her data ¹¹ .

meant to solve” (2004, p.85). Therefore, meeting all of the requirements tests below would be considered as a successful prototype for the purposes of this research.

To aid understanding, I have named the main test user Alice⁷. However, other users (intended to represent third parties such as the GP, hospitals, whānau, etc.) will be set up and will interact with Alice and her data. The final list of requirements is shown in table 5.1. Here, each requirement is given a unique identifier, and requirements are also grouped into functional categories. For examples, requirements RP1-3 are all discrete requirements which deal with the topic of creating and retrieving data. I will use these categories to group together development activity in the following sections.

It should also be remembered at this stage that, since the focus is on proving the concept that health data can be managed and shared in a distributed fashion, the above requirements are all very functional. That is, they describe basic functionality that the prototype should offer. In the real world there will be many other requirements – for example, utilisation of exchange

⁷It has been prevalent in cryptography to ‘humanise’ protagonists in this specific manner, ever since a 1978 paper on digital signatures (Rivest, Shamir, & Adleman, 1978). Rather than use Person A, Person B or Person C, for example, they used the names Alice, Bob and Carol. This practice has since become something of a phenomenon such that it has its own Wikipedia page (https://en.wikipedia.org/wiki/Alice_and_Bob), and has also been portrayed in the webcomic XKCD (<https://xkcd.com/1323/>).

⁸Referencing the first bullet point from section 5.1.4.

⁹Referencing section 5.1.1.

¹⁰Referencing the second bullet point from section 5.1.4.

¹¹Referencing section 5.1.2.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

standards for interoperability and identity management – but real decentralisation is such a novel concept that meeting the above requirements would already represent a meaningful contribution to the literature.

In summary we now have seven formal requirements that we can translate directly as tests, which can be used to assess efficacy of the prototype. In constructing an effective prototype we should expect each test to pass.

5.3. V1 Prototype design, build and evaluation by requirement category

This section will outline the design process for V1 of the prototype, including the holochain architecture and overview of functions and the security model for access control. This will also encompass design and results of the Holochain test suite (Tryorama), which are used initially to satisfy all of the requirements from the preceding section §5.2.

Firstly, I would like to clarify some key concepts and basic information about Holochain from a design and development perspective. Holochain is simply a *framework*. In programming, a development framework is a set of tools which permit the developer to save time. There may be some commonly utilised tasks that a framework can package and provide to the developer as a shortcut. This means the developer doesn't need to write every single piece of functionality entirely from scratch; they will be plugging in some of these shortcuts to their code, so that they can save time.

For example, frameworks might have macros or functions available for common tasks like converting dates, or they might offer novel functionality to help with common issues like caching data or making it easier to share data between functions. The developer still needs to do all the other routine development tasks (design and code), but it will be a quicker and easier experience for them. The Holochain development framework is called the 'Holochain Development Kit' (HDK), and its documentation notes:

“The HDK lets the developer focus on application logic and, as much as possible, forget about the underlying low-level implementation. It would be possible to write DNA source code without an HDK, but it would be extremely tedious!”¹²

In Holochain's case, development must be done in a programming language called *Rust*. There is a very good reason for this choice, and it is to do with exactly how holochain facilitates distribution.

¹²<https://docs.rs/hdk/0.0.46-alpha1/hdk/index.html>, although please note this is not the latest version of the HDK. That can be found here: <https://docs.rs/hdk/latest/hdk/index.html>.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

Since Holochain can offer true distribution, via peer-to-peer communication, it is very important that any Holochain application communicating over the internet can do so as efficiently as possible. One way of doing that, which is becoming both more mature and more prevalent, is to use a binary format called WebAssembly (usually abbreviated to ‘Wasm’). Wasm is now supported by default in the four main browser engines¹³, and is lauded as “a highly touted effort that not only is set to run web apps in the browser at near-native speeds but also allow for other languages to be used for browser programming beyond JavaScript” (Krill, 2017, para.5). Only a relatively small number of languages can currently compile to Wasm, however, and thus Rust was selected for the HDK.

Using Rust represented a particular challenge to myself, as a developer, since I had never used it before and it has both a very steep learning curve and an intimidating reputation (Yegulalp, 2018). It took a very intense learning period, with the help of an experienced mentor, to be able to produce working Holochain code and the final artifact from this DSR process.

Developing in the Holochain framework can be most easily understood in terms of three key components – data, functions and testing. I will discuss each of these in the context of all three requirement categories.

5.3.1. Recording and retrieval (RP1-3) – Data

Data storage in Holochain is analogous to that of Blockchain, in that the data is written to a ‘chain’ which is immutable. The notable difference is that Holochain is truly distributed since its network communication is peer-to-peer rather than via a central server or authority. The chain of data is stored on a users local device (private data), but can also be distributed amongst the entire network of users if appropriate (public data). I have made an important design decision early on to make the prototype data *private* only. This is discussed further in section 5.3.4.

Data can be conceptualised in two spaces within Holochain – the backend code (written in Rust), and within the Holochain application itself. In terms of backend code, all data objects must be written as a *struct*¹⁴. A struct is a datatype that houses other data types. For example, rather than having a single data type called **ConsultationNotes**, which must contain **text** (String), Holochain requires it to be written as a struct:

```
1 struct Consultation {  
2     ConsultationNotes: String  
3 }
```

Using this approach, the developer has to first access the object **Consultation** before they can access the data inside it¹⁵.

¹³<https://webassembly.org/roadmap/>.

¹⁴<https://doc.rust-lang.org/std/keyword.struct.html>.

¹⁵The rationale for use of structs in Holochain is because all data must be *serialisable*. That is, it must be stored in

5. Research Phases 2-3: V1 Prototype design, build and evaluation

The point of all this is to explain that one of the first development tasks is to understand what data needs to be captured, and in which format. In the case of using the HDK, one must think about the structs that are required and what will go inside them. When thinking about where to start, I decided to simply work through the list of requirements in order. The first requirement (RP1) was to implement the ability to record data to a users chain. The initial struct developed to achieve this was:

```
1 #[hdk_entry(id = "health_data", visibility = "private")]
2 #[serde(rename_all = "camelCase")]
3 #[derive(Clone)]
4 pub struct HealthData {
5     pub created_by: AgentPubKey,
6     pub content: String,
7     pub resource_type: String,
8 }
```

In this example we have a generic struct called **HealthData**, which contains three data points – who it was created by (**created_by**), what is the content (**content**) and what is the category of the data (**resource_type**). These items have different data types (two are **String** and one is **AgentPubKey** which is an HDK specific data type which permits unique identification of users in a network). The reader will also notice some additional lines at the top prefixed with ‘#’. These are Rust *attributes* which can be attached to an object and help do some of the heavy lifting. The first line, for example, is specifying that this data is a holochain entry (something that will live on a users chain), and that by default it is not made public.

The other space in which Holochain data can be viewed is where it has been processed by the backend code, and is available to the user within a Holochain app. Using the example above, we have defined the **HealthData** struct as an **hdk_entry**. This means it will be recorded as a Holochain entry in that users chain.

An *entry* is a specific concept in Holochain, and can be considered as the most basic unit of user data. But once the HDK recognises it as an entry, there is a process of adding it to the users chain and attaching important metadata. In this way, the HDK formally makes that data entry an immutable piece of data on a users chain, which is placed in order and can be searched and retrieved using other Holochain functions.

For V1 of the prototype, my goal was to quickly get to the point of storing and retrieving user data in Holochain. To do so, I needed to define some additional structs and the full list of these is shown below:

- **HealthDataInput**. Takes data inputted by the user (**content** and **resource_type** fields only) and stores in order to create a Holochain entry.

a format that can be transmitted and reconstructed reliably later in a different host system. A simple String value cannot achieve this, and thus all data must form part of a struct.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- **HealthData**. Picks up the data from **HealthDataInput** and dynamically adds a **created_by** field, based on the user making the entry. A Holochain entry is created from this struct.
- **HealthDataOutput**. Picks up all entry data from **HealthData**, and dynamically adds a unique **entry_hash** field which can be used as a unique identifier.

For the basic health data entry item, we now have three separate structs which can hold data. Now we need to find a way to connect these structs and do something useful with them; this is where functions are used.

5.3.2. Recording and retrieval (RP1-3) – Functions

A function is a piece of code that does a particular thing. For example, you may have a function which converts a temperature value from celsius to fahrenheit. The function will take one input (the temperature in celsius), it will then execute code which performs a calculation using your input data and finally returns the result in fahrenheit. The function does exactly one thing, and is very clear about what can be used as an input and what it will give you back.

For our RDHIS prototype, a relevant function might be to return all the health data for one user. The function structure might look something like this:

1. Send name of the current user to the function
2. Function scans the Holochain source chain, using some macros provided by the HDK, and bundles all the data together
3. The bundled data is then presented back to the user.

In fact this is exactly the logic of the function implemented to retrieve a users health data. as we will see below. The first version of the **create_health_data** function looked like this:

```
1 // Creates a new health_data entry
2 pub fn create_health_data(health_data_input: HealthDataInput,) -> ExternResult<
    HealthDataOutput> {
3     let agent_info = agent_info()?;
4     let health_data = HealthData {
5         created_by: AgentPubKey::from(agent_info.agent_latest_pubkey.clone()),
6         content: health_data_input.content,
7         resource_type: health_data_input.resource_type,
8     };
9     create_entry(&health_data)?;
10
11     let health_data_hash = hash_entry(&health_data)?;
12     let path = health_data_path();
```

5. Research Phases 2-3: V1 Prototype design, build and evaluation

```
13     path.ensure()?;
14     create_link(path.hash()?, health_data_hash.clone(), ())?;
15     Ok( HealthDataOutput{
16         entry_hash: EntryHash::from(health_data_hash),
17         entry: health_data,
18     })
19 }
```

Without needing to delve into everything in the above excerpt, it is worth pointing out the following key parts:

- Line 2. The function is given a name (**create_health_data**) and we specify that the input and output data must conform to the two named structs we have already defined.
- Line 3. We obtain the current user ID using an HDK function¹⁶.
- Lines 4-7. A new **health_data** variable is defined, which we conform to the **HealthData** struct and populate with values. The first **created_by** value is copied from the variable we made on Line 3, and conformed to the **AgentPubKey** data type. The other two struct items are simply copied from the input data using dot notation.
- Line 9. The **health_data** variable is used to actually create a Holochain entry on the user's chain.
- Lines 12-14. We define and create a path for the entry, to help us with easier retrieval if necessary. This utilises a **health_data_path()** helper function, which is not shown here.
- Lines 15-18. We define what a successful execution of the function will return to the client. In this case, we return all of the **health_data** variable used to create the entry, but add a hash of the whole entry to help us identify it uniquely. This value is conformed to the **EntryHash** HDK data type.

So the above function will statelessly be called whenever it receives an input, and will process all in the above manner. There is, of course, an additional function for *retrieval* of health data entries, and I will also show it here for completeness:

```
1 // Returns all health_data entries for the calling user.
2 // Requires no input and will just return a vector of every health data entry.
3 pub fn list_health_data() -> ExternResult<Vec<HealthDataOutput>> {
4     let path = health_data_path();
5     let links = get_links(path.hash()?, None)?;
6     links
```

¹⁶Users are referred to as Agents in the Holochain terminology.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

```
7         .into_inner()
8         .iter()
9         .map(|link| {
10             utils::try_get_and_convert::<HealthData>(HoloHash::from(link.
                target.clone()))
11             .map(|(entry_hash, entry)| HealthDataOutput {
                entry_hash, entry })
12         })
13         .collect()
14 }
```

This function is perhaps somewhat easier to understand for the layperson. It requires no user input (the user is simply saying, return all the data that belongs to me), and the output is potentially many instances of a **HealthDataOutput** struct which is formatted as a **Vector**¹⁷. Inside the function, the path that we set against all entries (**health_data_path()**) is used for retrieval and they are then collected¹⁸ into our **HealthDataOutput** vector which can be presented back to the user.

In summary, the following functions were all implemented for the V1 prototype in order to achieve writing and retrieval of Holochain data:

- **create_health_data**. Takes user input and records, together with the **AgentPubKey** identifier, to the users chain.
- **health_data_path**. Specifies a hard-coded path for all entries using this function as “health_data”. The concept with this is that we will probably require further types of entries in future, and this will aid retrieval.
- **list_health_data**. Requires no input, and will retrieve all entries on the calling users chain. Entries are collected as a vector and conformed to the **HealthDataOutput** struct, along with a unique identifier for that entry (**EntryHash**).

5.3.3. Recording and retrieval (RP1-3) – Testing / evaluation

Having structured the data, and designed functions that allow the data to be recorded, passed around and retrieved, the final piece of the puzzle is to check it all works using the Tryorama test suite. Tryorama is a tool provided by Holochain which allows the developer to test scenarios¹⁹.

For example if you have built a Holochain chess game, and you want to check it works before letting users access it, you would write some scenarios in Tryorama and check that all of them are passed successfully. Some scenarios that you may want to test are:

¹⁷Effectively just a growable list of objects – <https://doc.rust-lang.org/std/vec/struct.Vec.html>.

¹⁸<https://doc.rust-lang.org/std/iter/trait.Iterator.html#method.collect>.

¹⁹<https://github.com/holochain/tryorama#conceptual-overview>.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- Is the proposed move valid? (i.e. the Bishop only moves diagonally)
- Is it the correct players turn? (i.e. ensure a player cannot move twice in a row)
- If pieces have been taken, have they correctly been removed from play?

Tryorama will allow the developer to construct scenarios with players, and will then simulate a full run of your Holochain app which validates whether or not the tests you have specified have passed or failed.

In this way the Tryorama tool serves as a useful evaluation step within the DSR process. Hevner et al. note that “a design artifact is complete and effective when it satisfies the requirements and constraints of the problem it was meant to solve” (2004, p.85), and formal testing is one way that the evaluation criterion can be met. There are, of course, many ways of evaluating an artifact. Hevner et al. (2004) group these approaches into types, such as *Observational* or *Experimental*. In practice, the Tryorama evaluation I undertook against V1 of the prototype spanned across these²⁰. These approaches are summarised in table 5.2, where I have also added a final column to indicate its applicability to the test I propose to undertake in Tryorama.

We can see from this table that Tryorama occupies the Experimental and Testing categories of evaluation²¹. It will meet the criteria for a Simulation, because it allows a full run of a scenario with multiple active users interacting with a sandboxed version of a Holochain app. The methods shown under the Testing category are a little more complex.

The practical distinction between ‘Black box’ and ‘White box’ testing are somewhat less clear in 2022. Originally they presented two key paradigms in software testing – one that required no knowledge of the system itself but tested for expected functional outputs (Black box), and another that usually saw the developer (with deep knowledge of the system) build tests at low levels of the code to ensure that each part was working as expected (White box). ‘Gray box’ testing is a middle ground between these two, and suggests it is optimal to have understanding of the internal workings of the evaluated system to design the test cases – but to actually perform the testing at the ‘Black box’ level.

This approach is more fully aligned with working in Tryorama, where the developer must have knowledge of relevant backend functions and how they work in order to build and implement what may be viewed as the more ‘Black box’ component of simply assessing whether tests have passed or failed based on provided inputs. It should also be noted that typologies of artifact

²⁰In fact, many of these are not discrete evaluative activities and can easily overlap. For example a static analysis might be performed early in the development cycle, to find obvious errors in the code. This may be subsequently followed up by a functional test or simulation.

²¹Whilst ‘Dynamic analysis’ sounds applicable to my use case, it is really geared towards applications with a large number of possible outputs – as many as possible of which should be tested (Khatiwada, Tushev, & Mahmoud, 2018).

5. Research Phases 2-3: V1 Prototype design, build and evaluation

Table 5.2.: Design Evaluation Methods, adapted from Hevner et al. (2004).

Category	Method	Description	Utilised in Tryorama
Observational	Case study	In-depth real world study	No
	Field study	Monitor artifact use in multiple projects	No
Analytical	Static analysis	Examine artifact for static qualities or defects	No
	Architecture analysis	Examine how artifact fits into a technical architecture	No
	Optimisation	Demonstrate inherent optimal properties of artifact	No
	Dynamic analysis	Examine dynamic qualities of artifact while in use	No
Experimental	Controlled experiment	Examine artifact qualities in a controlled environment – traditionally achieved by holding all factors constant except one independent variable	No
	Simulation	Execute artifact with artificial data	Yes
Testing	Functional (Black box)	Test artifact to ensure that outputs are correct, given particular inputs – no knowledge of the internal workings are required	Partial
	Structural (White box)	Detailed source code testing, incorporating elements such as unit testing and integration testing	Yes
Descriptive	Informed argument	Utilise knowledge base to argue for artifact's utility	No
	Scenarios	Construct detailed scenarios to demonstrate artifact's utility	No

5. Research Phases 2-3: V1 Prototype design, build and evaluation

evaluation in the technology space (as shown by Hevner et al. in table 5.2) do not always line up easily with real world software testing approaches.

For example, ‘unit testing’ is an approach used by most programmers to ensure that what they are creating is meeting project requirements. In practice, a unit test may be carried out via static or dynamic analysis and, furthermore, the unit test is generally considered an important sub-method of the ‘White box’ method. This is all to say that evaluation and testing methods do not line up tidily for categorisation, and the evaluation methods presented by Hevner et al. (2004) may not have been set out specifically with app development in mind. Venable, Pries-Heje, and Baskerville go further and note that the influential Hevner paper in fact provides “no guidance on method selection or evaluation design” (2012, p.430). Additionally it is true that software testing approaches will, in any case, have advanced in the time since that paper was originally published.

One thing to note here is that the Rust programming language is very strict about what it permits. It is a compiled language, which means that the developer must compile all their code which then produces a binary file for execution. At the point of compilation Rust runs some very strict checks to ensure the code is syntactically correct. For example, Rust will not even compile successfully if a declared variable is unutilised. This is in stark contrast to other languages, such as Python, which will let you write almost anything you want. With these programming languages, errors will only surface at runtime. I point this out here to argue that Rust in fact has a type of in-built static analysis process via the compilation step. If it compiles successfully, you have functioning code with no errors, and this can be considered another type of automated evaluation activity.

Another consideration is that software testing takes different forms during the life cycle of an application. For example, individual unit tests may be written initially for each function to verify that they work as intended in isolation. Further tests then need to be developed to ensure that they work appropriately *in concert*. Any changes made as a result of these steps then need to undergo regression testing to ensure that, in fixing an identified error, we do not inadvertently introduce a new one. Essentially, any process of software development is in fact a continuous cycle of evaluation of different sorts²².

A further evaluative consideration within DSR is whether the evaluation activity is being performed *ex ante* or *ex post*. *Ex ante* evaluation is evaluation of a design or model – an “uninstantiated artifact” (Venable et al., 2012, p.430), whereas *ex post* evaluation is evaluation of an actual instantiation. Whilst Hevner’s seminal 2004 paper builds upon the work of March and Smith (1995) by suggesting that the DSR process is *brought to life* by the instantiation artifact²³,

²²Although a number of different frameworks have evolved – Software development life cycle (SDLC), Agile, Waterfall – all of them share a focus on repeated testing and evaluation at different levels of detail.

²³A successful ‘instantiation’ must “show that constructs, models or methods can be implemented in a working

5. Research Phases 2-3: V1 Prototype design, build and evaluation

there is no mention made of unit testing or continuous improvement process above and beyond the interplay between the evaluation and development phases.

I mention this because, at the point of having code which I believe will achieve requirements RP1-3, I technically have an instantiated artifact. Certainly, it only meets three of the requirements (as listed in section §5.2) and is far from complete. but it is nevertheless an instantiation. A working product has been created. This means that I will be conducting only *ex post* evaluation, since I do not start testing the meeting of requirements until working code has been implemented (or the instantiation artifact has been instantiated, albeit only partially).

In summary, I perform here an evaluation of the V1 prototype code to determine whether or not it meets the current requirements of interest (RP1-3, or the recording and retrieving of health data from a user's chain). This might be considered more of a *unit test* at this stage, since each action (recording and retrieval) is carried out by exactly one function. In DSR, however, the most important consideration is that the concept can be brought to life in a working system and verified as such by a structured process²⁴. Whilst this may initially only be a small test of a discrete component, we are still nevertheless testing for the working feasibility of distributed storage and retrieval of data. Its value should therefore not be underestimated, even at this early stage.

5.3.3.1. Defining a Tryorama test

With the above in mind, we simply need to define, build and implement the Tryorama test. Tryorama is a node.js package, and tests must be written in TypeScript²⁵. The workflow essentially involves:

1. Importing some Holochain libraries
2. Registering a scenario
3. Creating players (or users) who will take part in that scenario
4. Utilise the app functions you have built to simulate those players interacting in a network.

There is a moderate amount of 'boilerplate' code associated with this, which I won't replicate here²⁶. My test scenario involves setting up one player (Alice). When we move into more

system" (Hevner et al., 2004, p. 79).

²⁴The process must be *considered and structured*, but need not be entirely objective – as can be seen by the range of methods suggested in table 5.2.

²⁵A superset of JavaScript featuring static typing, which was developed by Microsoft and saw its first stable release in 2016.

²⁶A helpfully commented example test script is available as part of the Tryorama git repository at <https://github.com/holochain/tryorama#sample-usage>.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

complex functions, we will need to add in extra users – Bob and Carol. Tryorama will set each of these up as separate users, operating from their own sandboxed environments (using different port numbers on the host machine). Players will take actions as determined by the developer, by using the functions that have been developed in the Holochain app.

As discussed in section 5.3.2, we are going to utilise the two functions **create_health_data** and **list_health_data**:

```
1  const alice_result_1 = await alice.call(  
2      "radvue",  
3      "create_health_data",  
4      {  
5          content: "Alice's first entry",  
6          resourceType: "Primary care GP"  
7      }  
8  );  
9  
10 const alice_result_2 = await alice.call(  
11     "radvue",  
12     "create_health_data",  
13     {  
14         content: "Alice's second entry",  
15         resourceType: "Cardiology"  
16     }  
17 );  
18  
19 const alices_data = await alice.call(  
20     "radvue",  
21     "list_health_data",  
22     null  
23 );  
24  
25 t.equal(alices_data.length, 2, "We got exactly 2 entries back for Alice!");
```

I will describe what is happening here:

- Lines 1-8. Here we are calling the **create_health_data** function as the user Alice and recording some data. The **content** and **resource_type** fields contain some artificial data for test purposes. The result of this function call is saved as a constant **alice_result_1** so we can utilise it later, if necessary.
- Lines 10-17. We repeat that process again, but make a second entry with some different **content** and **resource_type** values.
- Lines 19-23. Here we are calling the **list_health_data** function as Alice, and passing the vector of responses to another constant **alices_data**.
- Line 25. This line is the actual test that we want to perform. It says that we expect there

5. Research Phases 2-3: V1 Prototype design, build and evaluation

to be exactly two entries recorded for Alice and, if so, to return a success message. If the length of `alices_data` is anything except two then the test will fail.

5.3.3.2. Executing a Tryorama test

We can now run the test by executing a simple command (`npm test`), where test is defined in our package.json file as "`npm run build:happ && npm t -w tests`". This command will compile and rebuild the app, and then execute the test file we have configured above.

This test passed the first Tryorama test. The full results are shown in figure 5.1. The output shows a few key things happening:

- Lines 14-18. The test script compiles the Holochain app and makes sure it is syntactically valid before running any tests.
- Lines 20-41. Basic configuration for the test is parsed and applied from the package.json configuration file.
- Lines 50-60. The player/user Alice is set up and their sandbox environment is activated.
- Line 61. The test is executed and evaluated. In this case, our success message which we configured above is printed on screen.
- Lines 63-65. The full test results are summarised. In this case, all (one) tests pass.

Having completed a design, development and evaluation/testing cycle for requirements RP1-3, I will now summarise the work done on the other requirement categories to complete the V1 prototype.

5.3.4. Access control (RP4-6) – Data

Moving into the requirement category of access control invites us to explore a key function of the Holochain framework. It could be argued that Holochain is primarily focused on public uses of data. That is, applications where data is shared amongst users. This is a real strength of Holochain, particularly in how the peer-to-peer model has significant benefits for data resilience and availability.

When data is made public in Holochain, each entry is shared randomly amongst all the users in the network. The result of this is that, individually, no user has access to all of the data but, collectively, everyone has access to all of the data. This can be better represented visually, as shown in figure 5.2.


5. Research Phases 2-3: V1 Prototype design, build and evaluation

Figure 5.1.: Tryorama test result for RP1-3

```
1 > test
2 > npm run build:happ && npm t -w tests
3
4
5 > build:happ
6 > npm run build:dnas && hc app pack ./workdir
7
8
9
10 > build:dnas
11 > npm run build:zomes && hc dna pack ./dna/workdir
12
13
14 > build:zomes
15 > CARGO_TARGET_DIR=target cargo build --release --target wasm32-unknown-unknown
16
17 Wrote bundle /data/holochain/radhis2/radhis/dna/workdir/radvue.dna
18 Wrote bundle /data/holochain/radhis2/radhis/workdir/radvue.happ
19
20 > tests@0.0.0 test
21 > TRYORAMA_LOG_LEVEL=info RUST_BACKTRACE=1 RUST_LOG=holochain::core::ribosome::host_fn::debug=debug TRYORAMA_HOLOCHAIN_PATH="holochain" node --loader ts-node/esm
--experimental-specifier-resolution=node src/index.ts | tap-diff
22
23 08:54:33 [tryorama] [32minfo][39m: Using the following settings from environment variables:
24 08:54:33 [tryorama] [32minfo][39m: {
25   "adminInterfaceId": "tryorama-interface-admin",
26   "appInterfaceId": "tryorama-interface-app",
27   "stateDumpOnError": true,
28   "zomeCallTimeoutMs": 90000,
29   "conductorTimeoutMs": 125000,
30   "strictConductorTimeout": false,
31   "chooseFreePort": false,
32   "logLevel": "info",
33   "portRange": [
34     33000,
35     34000
36   ],
37   "legacy": false,
38   "singletonAppId": "TRYORAMA_APP",
39   "holochainPath": "holochain",
40   "lairPath": "lair-keystore"
41 }
42
43 radvue tests
44   FIXME: ignoring onJoin
45   08:54:33 [tryorama] [32minfo][39m: Spawning lair for test with keystore at: /tmp/tmp.VrAinm0I9f/tryorama/xtjacX/keystore
46   08:54:33 [tryorama] [32minfo][39m: Using conductor path: holochain
47   08:54:33 [tryorama] [32minfo][39m: Holochain version: holochain 0.0.115
48   08:54:33 [tryorama] [32minfo][39m: Conductor config path: /tmp/tmp.VrAinm0I9f/tryorama/xtjacX/conductor-config.yml
49   08:54:34 [32minfo][39m:
50   @@@ [[CONDUCTOR c0]]
51   @
52   @ ###HOLOCHAIN_SETUP###
53   @ ###ADMIN_PORT:33000###
54   @ ###HOLOCHAIN_SETUP_END###
55   @ Conductor ready.
56   @
57   08:54:34 [tryorama] [32minfo][39m: Conductor 'c0' process spawning completed.
58   App Port spun up on port 38255
59   08:54:35 [tryorama] [32minfo][39m: conductor 'c0' exited with code null
60   FIXME: ignoring onLeave
61   ✓ We get exactly 2 entries back for Alice!
62
63 passed: 1 failed: 0 of 1 tests (3.7s)
64
65 All of 1 tests passed!
66
67
```

5. Research Phases 2-3: V1 Prototype design, build and evaluation

Figure 5.2.: Holochain Distributed Hash Table



In this visual representation we can see multiple users sharing data with a Holochain app. The coloured bars represent which users are an authority for specific entries or pieces of data (represented by the page images). It is clear to see that, together, the users have access to all the data but each individual entry resides only with two or three users. This is a good way of understanding the ‘Distributed Hash Table’ (DHT) concept that is a core part of Holochain.

The DHT is a “distributed database of all public data ... which is basically just a big key/value store”²⁷. It means that, for example, the yellow user shown in figure 5.2 does not have access to the three entries shown at the bottom of the image. But, using the DHT, they can pass a request clockwise through the green and cyan users in order to find out where the data resides and then access it. This offers the significant advantage of resilience and availability.

For example, if any user goes offline (by powering off their laptop or closing a mobile app) any data that they are storing or own becomes unavailable. Holochain is peer-to-peer, so there is no central authority or server to go and get the data from – a network consists only of active users who store data, and may give you permission to access it. Therefore, storing data against multiple users protects against this issue and means that there is generally always a copy of the data available to access (if you have permission). This works very well where the data is made *public*. As discussed, my use case is predominantly around access control of personal data which diverges from this principle. It demands that data is *private*; only the owner has access to it, by

²⁷https://developer.holochain.org/concepts/4_dht/#finding-peers-and-data-in-a-distributed-database.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

default.

In making this design decision I am therefore consciously discarding the many advantages of using a DHT. Readers will also note that – if use of public data in a DHT is a way to ensure availability or resilience when users are offline – then only storing data privately creates a real problem in this regard. If only one user has access to some data (their own health data) then, without shared public data, it will be impossible to access it should that user go offline. This represents a serious issue when thinking about a RDHIS.

There are certainly ways around this, however. For example, it would be possible to store encrypted data in a public DHT, and have access control determined by the sharing of public keys which would allow delegates to decrypt that data. If the user went offline, the encrypted data would still technically be available because it has been shared across the DHT – but only people with the correct permission (a decryption key from the owner) would actually be able to read and understand it. I have added this particular issue to the list of future research possibilities in section §8.3, since it certainly deserves further exploration. Within the confines of this thesis, however, I am electing to develop a prototype which is not using a public DHT as the primary means of sharing data.

5.3.4.1. How can private data be shared in Holochain?

Private data in Holochain, as we have discussed, resides only with its owner – it does not leave the users device in any way. Sharing data in Holochain is achieved via a concept known as *capabilities*.

Capability-based security is in fact an existing and well-known security model in computing²⁸. The concept is that a token of authority can be generated, which contains references to objects and associated access rights. The underlying assumption is that nothing has permission by default. Permission must be explicitly granted via issuance of a capability token which specifies exactly what can be accessed, how and by whom. In this way, we can be certain that all users or entities only ever have the exact capabilities (permissions) that they require (M. Miller, Yee, & Shapiro, 2003).

Holochain follows this approach and summarises in its documentation:

In order for others to call one of their functions, the callee first has to grant access to that function. They do this by writing a capability grant entry to their source chain that specifies the function name, the access level, and any optional information depending on the access level (a random capability token and/or a list of assignees). After that, Holochain will automatically check the credentials of any

²⁸There are many security models, and each has relative advantages and disadvantages for a particular use case.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

incoming function call to make sure they match an existing grant. When a grantor wants to revoke or modify access, they simply delete or update that grant entry²⁹.

The above quote clarifies that, when issuing capability tokens to Holochain users, what is actually happening is that permission is being given to execute functions *on the grantee's behalf*. So, for example, if Alice wanted to share her data with Bob the actual backend process is as follows:

1. Alice issues a capability grant stating that user Bob can call the `list_health_data` function on her behalf
2. This grant (consisting of Bob's public key, a random secret and the name of the function being given permission to) is stored on Alice's source chain
3. Alice shares the random secret with Bob
4. Bob finds the new grant on his source chain and claims a capability grant which is also written to his source chain
5. Bob can now execute the function that Alice has given permission to.

That is, Bob will be calling `list_health_data` as if he were Alice. This introduces some minor issues, which we'll look at in section 5.3.9.

Finally I should clarify that the requirements of a RDHIS also demand that two specific *types* of capability are implemented. Firstly, as we have discussed, Alice needs to permit other users to view her private health data. Secondly, however, a feasible RDHIS demands that trusted health professionals have the ability to actually *record* data on Alice's chain. This is a separate and quite distinct permission, and requires a different workflow.

5.3.4.2. Building data components for capabilities

What data components are required to implement capabilities? As we saw in section 5.3.1, it is important to plan out the data objects that will be required. I relied on two example sources³⁰ to help me understand the basic outline of a capability process. Following this review I determined the following data structs were required:

²⁹https://developer.holochain.org/concepts/8_calls_capabilities/#how-to-secure-functions-against-unauthorized-use.

³⁰One which is a basic guide from the main Holochain git repository (https://github.com/holochain/holochain/blob/develop/crates/test_utils/wasm/wasm_workspace/capability/src/lib.rs), and another which is a solution for a series of learning tasks designed by a Holochain community member (<https://github.com/holochain-gym/developer-exercises/blob/main/2.intermediate/5.capability-tokens/solution/zomes/exercise/src/lib.rs>).

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- **GrantCapAccess**. Stores two key items received from the grantee – **function** (the name of the function being given access to) and **agent** (the unique identifier of the agent who is being given the permission).
- **CapReceive**. A struct which houses the random secret and associated information which is sent to the grantee. Specifically, it contains **cap_secret** (the random secret) and **from_agent** (the unique identifier of the grantee).
- **HealthDataRemoteInput**. This is an extension of the previous HealthData-related structs we have built. In fact, it is exactly the same as the original **HealthDataInput** struct but we need to add an extra **filter_by_agent** field. This is used to designate the identity of the calling user (which can be checked for validity against the issued capability grant).

So for a basic implementation of capabilities, we now have two separate structs which specifically manage the capability workflow. We have also added a third struct, which extends the previous work and will enable the interaction with a user's health data via delegated permissions.

5.3.5. Access control (RP4-6) – Functions

The capability workflow is more complex than simply recording and retrieving data for a single user. It entails issuing grants, requesting tokens and then passing them through to a function to provide the correct authorisation – all without any central authority. Utilising capabilities we are in fact just repeating the same basic functionality seen in requirements RP1-3 – recording and retrieval of data. The difference is that users will be executing these functions *on behalf of other users*. I will not present any source code in this section, but it can of course be inspected in the git repository for this project³¹.

The functions that were developed to implement the capability workflow are:

- **create_capgrant**. Requires a populated **GrantCapAccess** struct as input. Uses this info to call specific Holochain functions which generate a secret (**generate_cap_secret**) and create a cap grant entry (**create_cap_grant**). The grant is sent directly to the grantee's source chain and, at the same time, a separate entry is made in the users chain recording the issuance of this capability.

³¹For the capability workflow specifically please see <https://gitlab.com/alexpoor/radhis/-/blob/main/dna/zomes/radvue/src/capgrant.rs> and https://gitlab.com/alexpoor/radhis/-/blob/main/dna/zomes/radvue/src/health_data.rs for the current versions of the functions referenced here.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- **call_list_cap_claim**. Requires only the grantee identifier as an input. This function then obtains a refreshed capability token, and uses it to call a child function **list_health_data_remote**. This function then uses the token, and the user identifier, to make a remote call to the grantee's chain using a hardcoded child function **list_health_data_with_capgrant**. This last function is essentially a replica of the original **list_health_data** function, which simply follows the designated path and aggregates all entries into a vector which is presented back to the calling user.
- **call_create_cap_claim**. This is probably the most complex function so far, since it requires the grantee to actually send health data along with all of the capability workflow. The function call requires a populated **HealthDataRemoteInput** struct as input. This again obtains a refreshed capability token, and then uses it together with the **HealthDataRemoteInput** data to call a child function **create_health_data_remote**. This function uses the capability token, user identifier, and the actual health data being recorded to make a remote call to the grantee's chain using a hardcoded child function **create_health_data_with_capgrant**. This is a replica of the original **create_health_data** function.
- **delete_cap_grant**. This is one of the simplest functions yet. It simply calls an existing Holochain function of the same name. The function takes only the unique identifier (**HeaderHash**) for the issued capability grant as an input. This immediately invalidates the issued capability grant and, while the user can theoretically continue to access it, it is no longer valid and will simply return an error.

It should be clear that adding the capability workflow generates a lot more code; there are actually eight discrete functions above which all play a part. I should note here that this could have been fewer, except there is a requirement that any functions exposed to the user can only take one input. Some of the interstitial workflow above (for example **create_health_data_remote**) is really just packaging up the incoming data in a way that the final function can accept and work with – this could theoretically have been done inside a single function, if the single input constraint was not in place.

5.3.6. Access control (RP4-6) – Testing / evaluation

In order to test the access control requirements, it is of course necessary to ensure that we have multiple users in our Tryorama scenario. So we will now introduce Bob and Carol, and set them up in our scenario. This is a simple configuration change at the top of our test script:

```
1 export default (orchestrator: Orchestrator<any>) =>
```

5. Research Phases 2-3: V1 Prototype design, build and evaluation

```
2   orchestrator.registerScenario("radvue tests", async (s, t) => {
3
4       const [alice_player, bob_player, carol_player]: Player[] = await s.players([
5           config, config, config]);
6
7       const [[alice_happ]] = await alice_player.installAgentsHapps(installation);
8       const [[bob_happ]] = await bob_player.installAgentsHapps(installation);
9       const [[carol_happ]] = await carol_player.installAgentsHapps(installation);
10
11       await s.shareAllNodes([alice_player, bob_player, carol_player]);
12
13       const alice = alice_happ.cells.find(cell => cell.cellRole.includes('/radvue.dna'))
14         as Cell;
15       const bob = bob_happ.cells.find(cell => cell.cellRole.includes('/radvue.dna'))
16         as Cell;
17       const carol = carol_happ.cells.find(cell => cell.cellRole.includes('/radvue.dna'))
18         as Cell;
```

This code is simply telling Tryorama to set up three separate sandbox environments for the additional two users. Line 10 specifically ensures that the users can all connect their environments together, otherwise they would not be able to access each others' data even with the appropriate permissions.

Now the scenario is set up we can start to build our scenarios and define our tests. In this round I am trying to test requirements RP4-6, which requires validation that:

- Alice can share data with specified users
- Only those users can access Alice's data
- Alice can revoke that access and the users will no longer be able to access the data.

We will run this scenario using the **call_list_cap_claim** workflow only. The scenarios and tests are defined in sections as follows:

```
1   let aliceCapGrant = await alice.call(
2       "radvue",
3       "create_capgrant",
4       {
5           function: "list_health_data_with_capgrant",
6           agent: bob_happ.agent,
7       }
8   );
9
10  t.ok(aliceCapGrant, "Alice has issued a valid capability grant to Bob");
11
12  let aliceGrantHeader = aliceCapGrant.header_hash;
```

Please note the key events from the above:

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- Line 3. Alice calls the **create_capgrant** function, and provides two input values.
- Lines 5-6. Both the name of the function being given access to, and the identifier of the grantee are used as inputs to the function.
- Line 10. We test the capability grant to ensure it is valid and the secret has been sent to Bob for claiming.
- Line 12. We store the unique identifier for this capability grant as a constant so we can use it later on.

```
1 let capGrantResults = await bob.call(  
2   "radvue",  
3   "call_list_cap_claim",  
4   alice_happ.agent  
5 );  
6  
7 console.log(capGrantResults);  
8 1  
9 t.equal(capGrantResults.length,2,"Bob got exactly 2 entries from Alice")
```

The key events from above are:

- Line 3. Bob uses the capability grant to execute the **call_list_cap_claim** function. The only required input is the identifier of the user whose function is being executed – Alice.
- Line 7. The results of this function call are printed so we can check them.
- Line 9. A test is executed requiring that exactly two entries are returned to Bob from Alice's source chain.

```
1 await alice.call(  
2   "radvue",  
3   "delete_cap_grant",  
4   aliceGrantHeader  
5 );  
6  
7 try {  
8 let capGrantResults2 = await bob.call(  
9   "radvue",  
10  "call_list_cap_claim",  
11  alice_happ.agent  
12 );  
13 console.log("Try to use the deleted cap grant:")  
14 console.log(capGrantResults2);  
15 } catch (err) {  
16 t.ok(err,"Bob cannot use revoked CapGrant")  
17 };
```

5. Research Phases 2-3: V1 Prototype design, build and evaluation

The key events from above are:

- Line 3. Alice executes the **delete_cap_grant** function. The only required input is the unique identifier for the issued capability grant (in this case, **aliceGrantHeader**).
- Line 10. Bob tries to execute the **call_list_cap_claim** function again anyway. It should fail because he no longer has permission.
- Line 15. This block executes if an error is thrown. An error should be thrown because Bob does not have permission to this function any longer.
- Line 16. A test is executed to check whether an error is thrown. If it is, the test is successful – confirming that Bob cannot execute Alice’s function any more.

Finally, we make one additional test to make sure that Carol – who has not been involved at all so far – cannot access Alice’s data:

```
1 try {
2   await carol.call(
3     "radvue",
4     "call_list_cap_claim",
5     alice_happ.agent
6   );
7 } catch (err) {
8   t.ok(err, "Carol cannot list Alice's data")
9 }
```

This code is almost identical to Bob’s last attempt to use the revoked capability grant. Carol tries to use the **call_list_cap_claim** function on Alice’s source chain but no capability grant has been issued, so it should fail.

At this point the Tryorama test results in 923 lines of output which is too much to show here. In addition to the test we wrote in section 5.3.3, we now have an additional four. An initial run of the test script shows that all five tests have passed.

5.3.7. Audit (RP7) – Data

Requirement RP7 demands that each user has a full audit trail of both the issuance of any permissions (capability grants) and the actual utilisation of those capability grants by others. Holochain does not provide any mechanism to achieve this³² and so it was built from scratch, and represents a novel addition to generic Holochain apps. I would like to briefly flesh out what is required in terms of data for both these workflows.

³²Please note that the documentation at https://docs.rs/hdk/latest/hdk/capability/fn.create_cap_grant.html states that “Capability grants are explicit entries in the local source chain”, implying that they can be retrieved like any other entry. I could not discover exactly how to make this work, however, and so I decided to build a separate set of entries that I knew I could reliably retrieve.

5.3.7.1. Logging permissions

In this step we need to make a record of any capability grant that has been issued. The record will live on the users source chain, and therefore needs to incorporate the following data points:

- **Grantee.** The user who is being granted access.
- **Function.** The specific function which the grantee is being given access to.
- **Identifier.** A unique identifier which can be retrieved and used for use with the `delete_cap_grant` function.
- **Datetime.** A timestamp recording when the capability grant was issued.

The above data is captured in a **CapRecord** struct. The intention is that issuance of a capability grant will also execute some code that will populate this struct and store the entry. This is done as part of the `create_capgrant` function execution, which creates the **CapRecord** entry and assigns a specific path (`caprecord`) so that retrieval is quicker and easier. The entry and retrieval of this data is identical to that of `create_health_data` and `list_health_data`.

5.3.7.2. Logging capability grant utilisation

This workflow is also analogous to that of `create_health_data` and `list_health_data`. We are creating a specific set of entries and ensuring they can be retrieved. Because there are two separate capability workflows (list or create), there are also two routes for data to arrive into this struct and each will record separate values. With that in mind, the following structs were designed for use:

- **AuditData.** This struct is essentially the identifier for the calling user (the grantee) and a timestamp indicating when the request was made. There is also a `record_header` field, which will clearly delineate between create or list actions, as well as a `resource_type` field, which is only relevant for the create route.
- **AuditDataOutput.** Picks up all entry data from **AuditData**, and dynamically adds a unique `entry_hash` field which can be used as a unique identifier.

5.3.8. Audit (RP7) – Functions

The audit functionality can be thought of as a byproduct of capability grant use. It is not something that a user explicitly engages in, but is triggered by other actions. Therefore, much of the audit workflow takes part within existing functions. For example, both access control (RP4-6)

5. Research Phases 2-3: V1 Prototype design, build and evaluation

workflows incorporate logic which builds our audit layer. Within both **list_health_data_with_capgrant** and **create_health_data_with_capgrant** there is a block which writes audit data entries on behalf of the grantor. For example, the following is found in the middle of the **list_health_data_with_capgrant** function:

```
1 // ### write audit data
2 let now = sys_time()?;
3
4 let audit_data_input = audit_data::AuditData {
5     created_by: pubkey,
6     record_header: "LIST DATA".to_string(),
7     resource_type: "Pending".to_string(),
8     timestamp: now.to_string(),
9 };
10
11 create_entry(&audit_data_input)?;
12 let audit_data_hash = hash_entry(&audit_data_input)?;
13 let path = audit_data::audit_data_path(AgentPubKey::from(agent_info)?.
14     agent_latest_pubkey.clone());
15 path.ensure()?;
16 create_link(path.hash()?, audit_data_wrapper_hash.clone(), ());
17 // ### finish writing audit data
```

This function is logically identical to **create_health_data**. We are populating the input struct **AuditData** with required values. The user identifier is passed to the **created_by** field, the current time is passed to the **timestamp** field and the other two values are hard coded. This means that in a list potentially full of audit entries, a grantee can quickly identify who has been using which permissions and when. Line 13 creates a specific path for these entries, so that retrieval is quicker and easier.

As the reader may by now imagine, the process for retrieving these audit entries is logically identical to **list_health_data** and is predictably called **list_audit_data**. The function requires no input, it simply retrieves all entries from the **audit_data** path, aggregates them into a vector and presents the results back to the requestor. There is therefore only one dedicated function for audit entries but, as we have seen, the creation of audit data is mainly embedded into other functions.

5.3.9. Audit (RP7) – Testing / evaluation

Our Tryorama test so far has been extended to incorporate the issuing and usage of capability grants. As we have seen in the preceding section, this workflow should already have generated audit data for us to review; we simply need to modify the scenario to capture that information and build a test to check it. To thoroughly ensure the audit layer is working properly, the following tests will need to be implemented:

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- Check that Alice has 0 CapRecord or AuditData entries at the beginning of the scenario (none have been issued and none have been utilised).
- Issuance of a capability grant must appropriately record CapRecord data. Alice can confirm there is 1 CapRecord entry.
- Bob's use of the capability grant appropriately records AuditData data. Alice can confirm there is 1 AuditData entry.
- The grantee can review any of this data to confirm accuracy.

I will now add the following components to the Tryorama test script:

```
1 let aliceAuditCheck1 = await alice.call(  
2     "radvue",  
3     "list_audit_data",  
4     null  
5 );  
6 t.equal(aliceAuditCheck1.length, 0, "Alice has 0 Audit Data entries");  
7  
8 let aliceCapCheck1 = await alice.call(  
9     "radvue",  
10    "list_caprecord",  
11    null  
12 );  
13 t.equal(aliceCapCheck1.length, 0, "Alice has 0 Cap Record entries");
```

These two tests are run at the beginning of the test script. They are simply checking for the presence of any **CapRecord** or **AuditData** entries. Since we have not yet issued any capability grants in our scenario, these should both be empty and return 0 results. The same tests can be run later, where we would expect to find one **CapRecord** entry (Alice has issued one capability grant to Bob), and one **AuditData** entry (Bob has successfully used the capability grant only once).

At this stage all the tests (there are now nine in total) passed. However, we have not yet actually returned any of the entries for checking. It was at this point that I noticed an issue with the **AuditData** entries. Because of some misconfiguration in the flow of data between functions, all the **AuditData** entries were being recorded with Alice's identifier. This was not appropriate because **AuditData** is supposed to show the grantee who accessed (or wrote) the data. I therefore had to backtrack and make sure that the identifier for the grantee was properly passed through the workflow so it could be recorded inside **AuditData**.

This was an interesting issue because it could easily have been masked by the test results, without either manual review of the returned data or a specific test built to verify the correct identifier was present in the entry. In fact, I did attempt to build a specific test that would ascertain Bob's

5. Research Phases 2-3: V1 Prototype design, build and evaluation

unique identifier and check that it matches the `createdBy` field within `AuditDataOutput`. Despite being able to verify visually that the two values were indeed the same – and thus expecting that the test passed – the test kept failing.

I determined that this is likely to be a data type issue, since these identifier values are stored as a ‘Buffer’ which is “raw memory allocation outside the V8 heap”³³. There are certainly methods by which to convert Buffer data to other data types, but this was not straight forward in TypeScript and I deprioritised this after *manually* verifying that this particular test had passed.

5.4. Summary

Overall the first round of design, build and evaluation went very well. There were certainly many hours of frustration and confusion, but slowly building my knowledge with the expertise of my Holochain mentor finally resulted in a working RDHIS prototype. It was doing everything I originally set out to achieve, and was truly distributed.

According to the commit history, available for review in the project git repository, the final commit was made on 21 November 2021 and represents all the development work on the V1 prototype up to that point. A total of 131 commits were made to the repository, encompassing a total of 1,204 lines of code. A snapshot of the full codebase at this juncture can be found at <https://gitlab.com/alexpoor/radhis/-/tree/38ab9aabb5c7cef5d0c30b11a0c3334962ff14f5/>. Having built and passed all the tests, which were themselves derived from the gathered requirements, I was optimistic that the project may be nearing completion.

However, three factors changed my perspective on this:

- Around the same time, I was organising interviews for research phase four. In discussing my topic with potential participants, they were very interested but were especially curious to see and interact with the ‘product’. Of course, there was no such product to interact with. The only evidence of my hard work was an obscure test suite, which could only be run by a user who had already installed all the prerequisites to build a Holochain development environment *and* had cloned my git repository. In short, the barriers to access for anyone wanting to test the prototype were unacceptably high.
- Whilst my supervision team were very supportive of my progress, they too fed back that it felt rather intangible and that it may have less integrity throughout examination if it was felt to be a ‘black box’ process which no one else can really engage with and understand.

³³https://microsoft.github.io/PowerBI-JavaScript/classes/_node_modules__types_node_globals__d_.buffer.html.

5. Research Phases 2-3: V1 Prototype design, build and evaluation

- I was motivated to share my findings and particularly wanted to be able to carry out live demonstrations of the app to interested parties, for example at conferences. As things currently stood, it was not presentable except in the abstract and, for such an abstract topic as distribution of data, this represented a real obstacle to dissemination of my findings.

Whilst the above seems to sit outside the formalised rigour of an evaluation process, it is interesting to note that both Venable et al. (2012) and Gazem et al. (2018) consider ‘Artificial Evaluation’ one of two main categories of evaluation (the other being ‘Naturalistic’). Artificial Evaluation can take different forms, but relies on the feedback of expert users – in my case, this incorporated my supervision team, and the experts I was proposing to interview for research phase four. The clear feedback from this informal Artificial Evaluation was that a user interface was required.

Taking the above into consideration, I felt that the only path forward was to build a user interface for the code I had developed so that the app could be shared and used by anyone. I was especially anxious about this, as I am not by any means experienced with frontend development, but I knew enough to recognise that it would be a very large amount of additional work. My final commit on the project, after completing the user interface, was made on 19 February 2022. This clarifies that it did take an additional three months of focused work – along with an additional 163 commits and an additional 2,843 lines of code – to build this user interface. I will discuss this process in the context of a V2 prototype in the following chapter.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

From the outset, it had not been my intention to build any user interface¹. My concern was that, in building a user interface, the avenue of inquiry is widened much further into areas such as human-computer interaction or user-centred design – all of which I know nothing about, and would not have the space to explore properly in this thesis. Indeed, Hevner et al. specifically note that, depending on the artifact, there must be evaluative criteria for *design* and *style* – “design evaluation should include an assessment of the artifact’s style. The measurement of style lies in the realm of human perception and taste” (2004, p.86). Whether I intended to or not, I am producing an app with a user interface. Does this now require me to formally evaluate design and style, even if the user interface is not the artifact I am interested in from a research perspective?

My concern about this centred around the supposition that building a user interface simply *introduces* scope for a whole new field of critique around design, layout and usability. This is because it is practically the only thing that a user can engage with. All of the interesting distribution work is still done behind the scenes; it is not something tangible that an average user can measure. In some respects, it posed the risk of distracting from the core message around distribution. Yet I did also recognise the opportunity to aid dissemination of the findings – a very important thing.

I must also note that, around the time of this decision, the Holochain community had published some very useful tools which were designed to both accelerate app development and app sharing. The first is a tool called ‘scaffolding’, which asks the user some questions about what they want to do and then produces a complete file structure for a new Holochain app along with some example code². The second tool is called ‘launcher’, which is a small application that can be installed on any operating system³. The user can then install Holochain apps *into* launcher, and run them natively from their own device. Both these tools were significant developments within the Holochain community, with the former being an aid for developers and the latter making it

¹In fact I had specifically ruled it out in section §3.2.

²<https://github.com/holochain/scaffolding#rad-scaffolding-tools-for-holochain-applications>.

³<https://github.com/holochain/launcher#holochain-launcher>.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

much easier to share and use Holochain apps.

6.1. V2 Prototype requirements

All the requirements from table 5.1 were met in the V1 prototype. I am now adding the two below requirements (RP8 and RP9) to reflect the informal Artificial Evaluation also carried out as part of that:

Table 6.1.: V2 Prototype requirements

Category	Requirement ID	Requirement	Status
Recording and retrieval	RP1	A range of health data can be posted to Alice's record.	Complete
	RP2	Only Alice has read access to this data by default.	Complete
	RP3	Alice can review the data being posted to her record.	Complete
Access control	RP4	Alice can share her data with any other stakeholder, at any time.	Complete
	RP5	Third parties with delegated access can review the shared data.	Complete
	RP6	Alice can reliably edit or revoke the delegated access at any time.	Complete
Audit	RP7	Alice can review audit data, showing delegated permissions and details about third party use of her data.	Complete
Accessibility	RP8	The app has a user interface which anyone would be able to interact with for trial purposes.	-
	RP9	Users can easily install and test the app themselves.	-

The new requirements have been grouped under 'Accessibility'. Although this term actually means something rather specific in computing⁴, it was fitting here because it speaks about how an

⁴It certainly overlaps with usability, but has a particular focus on people with disabilities and aims to ensure that all users have equal access to digital products. My use of the term here certainly speaks about equal access, but more from a developer vs end user perspective. I certainly don't wish to minimise or dilute the importance of Accessibility specifically for people with disabilities.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

average user might be able to pick up the prototype Holochain app and use it because, otherwise, Holochain is too difficult to use.

With regard to RP8, however, I would add a caveat that the specific usability of design or style of the user interface will not be evaluated independently. I simply present an interface so that people can test and engage with the app themselves, or with a group – one outcome being to verify that none of the V1 prototype tests discussed in chapter 5 were hard-coded.

6.2. User interface design, build and evaluation (RP8)

Building a user interface for some already functioning backend code is not straight forward. When building an app where a user interface is a requirement (as with most apps), the interplay between backend and frontend code will significantly affect how each is designed and implemented. In my case, I only ever intended to write the backend code and, by this stage, have a fully functional prototype. Adding a user interface will necessarily require unravelling some of that work, and adding new components for the frontend to utilise. I will describe this process in this section, whilst also discussing the choice of toolset, and my design approach.

6.2.1. Choice of framework

In section §5.3 we discussed Holochain as a framework. Holochain was the most promising framework I could find which met the requirements for this project. When it comes to frontend frameworks, however, there are many to choose from⁵ and each has its own particular way of doing things. Since I had never used a frontend framework before, there was no *default* selection I could make based on familiarity.

An additional consideration was – *which frameworks are known to work well with Holochain?* Given my lack of knowledge I was largely dependent on asking Holochain community members for advice, and reading about other people’s experiences. One helpful constraint was that the software library which allows the connecting of a frontend to the Holochain backend was only available as a node.js package⁶. This meant that the frontend framework would have to be JavaScript based (which it most likely would have been in any case).

In selecting a framework, I found the Holochain ‘scaffold’ tool very useful indeed. As already mentioned, the tool will ask for some parameters and then create a working folder structure, complete with some example code. The tool has options to utilise one of three built-in frontend

⁵The Wikipedia page for [Comparison of JavaScript-based web frameworks](#) lists 21 frameworks, although there are indeed many more, and the landscape obviously changes rapidly.

⁶<https://github.com/holochain/holochain-client-js#holochain-client---javascript>.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

frameworks: Svelte⁷, Lit⁸ and Vue⁹. Of these, both Svelte and Vue are also noted by Kinsbruner and Bahmutov (2022) as being in the top five of the most utilised frameworks in 2022. I experimented with each. While Svelte, particularly, is rapidly gaining popularity – and was voted the most loved web framework by developers in the 2021 StackOverflow developer survey¹⁰ – I could not easily amend the supplied sample files to work with my existing code. I experienced the same with Lit. Vue, however, seemed to me to be very easy to adapt and amend to work with my code and I was rapidly able to get some basic functionality working. I’m certainly not saying that Vue is better or easier than the others; simply that I personally found it easier to work with initially and I decided to follow that momentum since time was of the essence.

I should also note that Vue is the most mature of these three and has a very large documentation and user base, which reassured me that I would be able to find answers to questions as I moved through this development process. My starting point was to tinker with the scaffolded Vue output and start to connect to some of my Holochain functions. I supplemented this with online resources; Tania Rascia’s guide on creating and updating data was particularly useful¹¹.

6.2.2. Frontend design

As I have already mentioned, it was not my intention to write a thesis on user interfaces or user-centred design. Nevertheless, I certainly needed to make the frontend functional and as usable as reasonably possible. Recognising this was outside my knowledge base, I consulted with an expert who works as a web developer¹². We discussed some layout options and came up with the below wireframe diagrams.

6.2.2.1. Landing/login page

Upon loading the app, the user should ‘land’ on a login page. I had not considered the issue of identity management, or user identity, before. Holochain identifies each user by assigning them a unique hash. As noted in section 5.3.9 this takes the form of a Buffer object, and appears as:

```
1 <Buffer 84 20 24 b0 47 a8 ac 42 49 99 57 9f a9 18 43 18 d0 6a 67 bf 38 31 9d f0 e1 4c
   b0 dd a8 04 7b a1 dd 77 ec 72 9e 28 7b>
```

.

⁷<https://svelte.dev/>.

⁸<https://lit.dev/>.

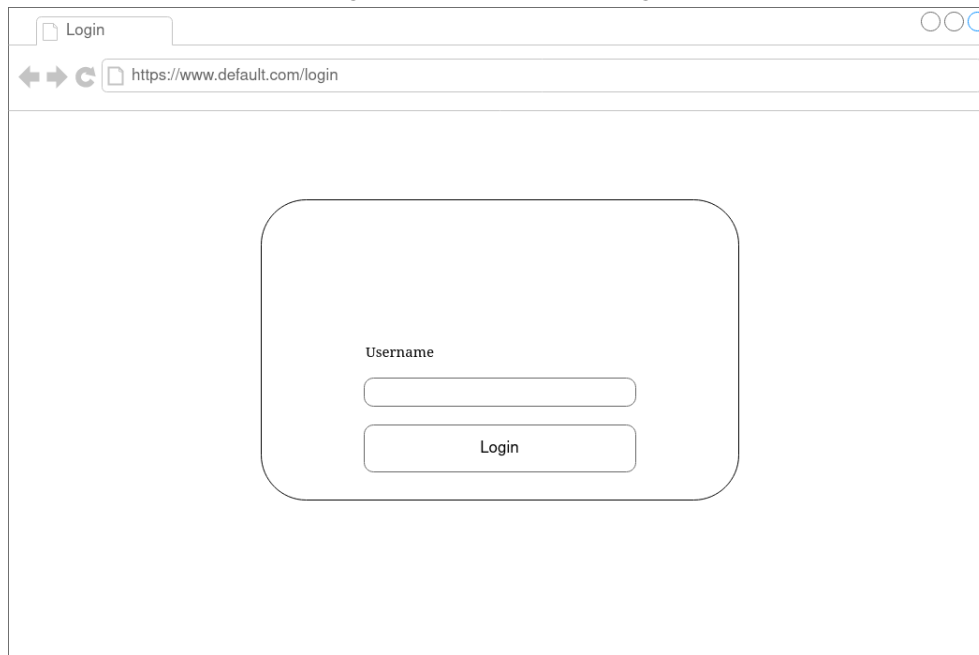
⁹<https://vuejs.org/>.

¹⁰<https://insights.stackoverflow.com/survey/2021#section-most-loved-dreaded-and-wanted-web-frameworks>.

¹¹<https://www.taniarascia.com/getting-started-with-vue/>.

¹²Lance Haysom, of polycode - <https://www.polycode.co.nz/> (as at May 2022).

Figure 6.1.: Wireframe: Login



These are utilised heavily in my backend code but, for obvious reasons, are not usable in the frontend. Because key requirements of the prototype are that users can share data, it is necessary for them to know who each person is. This therefore requires a mapping of Holochain identifiers to ‘usernames’, so that the frontend can present user-friendly names. This is the rationale behind the login screen wireframe, shown in figure 6.1.

On this page, the user simply enters a username and code in the backend will then need to map this to their Buffer object so that users reliably know which user is which. Having completed this step, they are then permitted into the app proper.

6.2.2.2. Main/data page

After the username has been submitted, the user will be directed to the main data page shown in figure 6.2. This page has two main panels – one for creating data, and another for viewing it.

The reader will recognise the **HealthData** struct fields shown in the ‘create data’ panel. These are a direct map to the data required in the backend, except two additional datetime fields have been added. On the right side is another panel which will simply show any **HealthData** for that user – whether they recorded it, or someone else did. This panel is simply calling the **list_health_data** function for that user, and presenting the Vector of results as a table.

Upon discussing how these panels would work, I reasoned that the review data panel could

Figure 6.2.: Wireframe: Main/data

The wireframe shows a web browser window titled 'My data' with the URL 'https://www.default.com/data'. Below the browser window, there are two main sections: 'Create data' and 'Review data'. The 'Create data' section contains four input fields labeled 'Content', 'Resource type', 'Start', and 'End', followed by a green 'Submit' button. The 'Review data' section contains a table with three columns: 'Content', 'Resource', and 'Date'.

Content	Resource	Date

become difficult to use if lots of data was being shown there. There may not be enough data points to sensibly distinguish between entries. This was the rationale for adding start and end datetimes to the form.

6.2.2.3. Access page

The Access page is intended to be the place where anything to do with delegation or sharing of data is performed. This meant that it is probably the busiest page, and potentially the most complex for users. The original wireframe is shown in figure 6.3.

The 'Grant access' panel on the left is where users can delegate permissions. They simply select the name of any other existing user in the network (this list will refresh whenever a new user joins), choose the desired permission (list or create) and then submit. In the backend this executes either the `call_list_cap_claim` or `call_create_cap_claim` functions for that user in order to create a capability grant.

The other two panels on this page are intended for users who have been given a capability grant. If you have been granted the 'List' permission for another user, you can select their name in the middle panel and click submit to see all their health data entries in a pop-up window. If you have been granted the 'Create' permission then, again, you just select that user in the rightmost panel and then complete the rest of the form. The data you enter is posted to that user's source

6. Research Phases 2-3: V2 Prototype design, build and evaluation

Figure 6.3.: Wireframe: Access

The wireframe shows a web browser window titled 'Access' with the URL 'https://www.default.com/access'. Below the browser window is a navigation bar with three tabs: 'Data', 'Access' (which is highlighted), and 'Audit'. The main content area is divided into three rounded rectangular panels. The first panel, 'Grant access', contains a 'Pick a user' input field, a 'Permission type' input field, and a green 'Submit' button. The second panel, 'List user data', contains a 'Pick a user' input field and a green 'Submit' button. The third panel, 'Create user data', contains a 'Pick a user' input field, a 'Content' input field, a 'Resource type' input field, a 'Start' input field, an 'End' input field, and a green 'Submit' button.

chain and audit data should also be available as a result of your use of their capability grant.

6.2.2.4. Audit page

Finally, the Audit page is the place for users to review information about permissions they have delegated. The wireframe is shown in figure 6.4. This is primarily a read-only page, since the information should just appear here whenever it's created. On the left, the user will be able to see any uses of an issued capability grant. For example, every time a grantee reviews the grantor's data a new entry with the date and time of access will be recorded here.

On the right is a panel which lists all capability grants that have been issued by that user. It will show who they were issued to, what permission was given and the date of issue. The exception to the read-only focus of this page is the 'Active' column on the far right. This is where a user can revoke a capability grant. All grants will show in this panel, and the user can simply pick which one they want to revoke and click an icon to do so. The revoked grant will still show in this table, for audit purposes, but it will be flagged as inactive.

6.2.3. Ex ante evaluation

Unlike the evaluation process described in chapter 5, discussing the design phase with an expert and being able to draft up some diagrams provided opportunity for *ex ante* evaluation. That is, a

6. Research Phases 2-3: V2 Prototype design, build and evaluation

Figure 6.4.: Wireframe: Audit



review of the wireframe diagrams could be conducted prior to any code being written.

What I found, however, was that requirement RP8 was sufficiently high-level enough to make this process difficult. My requirement for frontend development is “to make a working user interface”. In actual fact, my practical requirements are to ensure that all prior requirements (RP1-7) can be achieved via the user interface. This makes the *ex ante* evaluation process somewhat clearer and, indeed, a number of issues and problems were identified as a result of this process. I have recorded the outcomes of this process in table 6.2 as they relate to each requirement,

There were many other design-related issues that fell out of this *ex ante* evaluation (for example, *how can a user be sure their action has been successful?*), but these are not shown here because of our focus on the original prototype requirements.

The reader will notice that several of the evaluation findings relate to dynamically updating a page, or data, when another user has taken some action. For example, if you are looking at the Audit page then a grantee’s use of your capability grant should appear on the page without your having to take any action. This is indeed technically possible, but requires use of more nascent Holochain functionality which I did not have time to explore¹³.

The remaining items all require development work of some kind:

- EA1. Add two timestamp fields to the **HealthData** struct.

¹³https://developer.holochain.org/concepts/9_signals/#local-and-remote-signals.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

Table 6.2.: Ex ante evaluation findings

Finding ID	Requirement ID	Issue	Solution
EA1	RP3	Difficult to distinguish amongst large number of entries	Add start and end fields to HealthData struct
EA2	RP3	Entries are not updated automatically when made by third parties	Implement refresh button
EA3	RP4	Require user-friendly names to identify people	Implement login step with new Holochain Profile entry type
EA4	RP4	User list does not refresh independently	Implement refresh button for user list
EA5	RP5	Grantees do not know they are grantees	No push notifications available, park for future development
EA6	RP5	User list does not refresh independently	Implement refresh button for user list
EA7	RP6	User needs to easily see which grants are active or inactive	Implement a new Holochain CapRecord entry type, where the active field can be maintained
EA8	RP6	App does not have access to identifying HeaderHash for use with delete_cap_grant	Add HeaderHash to new Holochain CapRecord entry for easy retrieval
EA9	RP6	Grantees do not know grant has been revoked	No push notifications available, park for future development
EA10	RP7	Grant usage list does not refresh independently	No push notifications available, park for future development

6. Research Phases 2-3: V2 Prototype design, build and evaluation

- EA2. Frontend to add refresh button.
- EA3. Creation of new **Profile** entry type, and associated backend functions to create and retrieve **Profile** data.
- EA4. Frontend to add refresh button.
- EA6. Frontend to add refresh button.
- EA7. New **CapRecord** entry type to store active flag against all issued capability grants.
- EA8. New **CapRecord** entry type to store capability grant **HeaderHash** value so it can be used for deletion when required.

6.2.4. Frontend development and ex post evaluation

Having clarified the specific development tasks that had to be completed in order to meet requirements, I now began work on building the user interface. As already noted, this took three months to complete and required a large amount of additional code to be written.

As for testing, I should note that there are a wide range of frontend testing tools and approaches – the fact that entire *books* have been written on the topic (for example Kinsbruner and Bahmutov (2022) or Mwaura (2021)) indicate that it is a discipline in its own right. Formal test processes for frontend development do require the setup of software, and the generation of test scripts which can automate testing. This would be extremely useful in a bigger project but, again, I did not have the time or space within the confines of this thesis to learn another discipline entirely from scratch.

This is not to say that no testing or evaluation was performed. As with the V1 prototype, I tested along the way – ensuring that each workflow was operating correctly before moving onto the next. This is equivalent to a form of *unit testing*, which is a test to make sure that the individual components are working as intended. As we have seen in section 5.3.3, unit tests may be considered a subset of the ‘White box’ method, as categorised by Hevner et al. (2004).

While Kinsbruner and Bahmutov (2022) specifically note that it doesn’t matter whether testing is manual or automated, they do point out key functional and non-functional areas that should be tested and these are summarised in table 6.3.

Here we see testing dimensions which are focused on best practice for *production*¹⁴ applications. They are certainly excellent principles to consider in such a scenario. They are not all applicable for the V2 prototype, however, because of certain boundaries I have already imposed

¹⁴Production is the final stage of a development process, and indicates an application that is live and being utilised by real users.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

Table 6.3.: Testing dimensions for web apps, adapted from Kinsbruner and Bahmutov (2022).

Category	Dimension	Description	Relevant
Functional	Links	All links are functional, and direct user to intended destination.	Yes
	Forms	Mandatory fields are configured, and any required data types are enforced.	Yes
	Cookies	Check implementation of cookie policy, if applicable.	No
	Localisation	Application works equally well in different languages, if applicable.	No
	Responsiveness	Application works effectively on small screen resolutions (e.g. cellphones).	No
	Usability / UX	User experience has been checked.	Yes
Non-functional	Security	Verify restriction of access to secure content, require authentication where applicable, utilise SSL.	Partially
	Performance	Application is responsive with low latency, and stress testing has identified breaking points.	No
	Accessibility	Application is accessible to people with diverse abilities or disabilities.	No

6. Research Phases 2-3: V2 Prototype design, build and evaluation

in this thesis. I have marked each dimension as such in the table. The reasons for deeming these dimensions out of scope are as follows:

- **Cookies.** The application does not utilise cookies at all, in its current form.
- **Localisation.** I have only considered use of Aotearoa New Zealand English in the prototype. Aotearoa New Zealand has two written official languages and, of course, an ethnically diverse population. Real world usage of this app would require localisation into the most commonly used languages¹⁵.
- **Responsiveness.** At the time of writing, it is not possible to publish Holochain apps as mobile apps. The ‘launcher tool’, used to distribute and load Holochain apps currently works only on desktop computers. With that in mind, I did not consider device responsiveness as part of the prototype.
- **Security.** I have marked this ‘Partially’ merely to separate out the internal app security from more general IT security concerns. For example, the example security issue noted around access to secure content is in fact a core requirement of the prototype (RP2, RP6 and RP7). However, more general security concerns are out of the control of the author at present, and are dependent on upstream work being done on the Holochain framework. For example the ‘launcher’ tool does not implement encryption at rest currently, which is not best practice – particularly when storing sensitive health information.
- **Performance.** Under a ‘best practice’ scenario, an application would be rolled out across a fleet of servers¹⁶ that are specified to deliver high performance at all times. This prototype is designed only for use with an individual’s device – as a proof of concept it would be expected that there may be performance issues from time to time and no performance guarantees are made.
- **Accessibility.** Again, as a proof of concept prototype, accessibility was deemed out of scope. This is not to say that I considered it unimportant. I simply reasoned that I could get it to the point where it could be verified as proving the concept without yet spending additional time to make it accessible in the way that one would of a real *product*¹⁷.

This left the dimensions of Links, Forms, Usability/UX and Security to evaluate.

¹⁵<https://covid19.govt.nz/> is an exemplar of web content that has been localised to other languages, with some 28 localised translations available.

¹⁶Or an auto-scaling server pool, and/or use of a Content Delivery Network (CDN).

¹⁷I would also like to acknowledge here the same site <https://covid19.govt.nz/> as being an exemplar of accessible design, with provision for Braille Ready Format (BRF) and large print.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

Links and Forms are straightforward unit test tasks, and were checked by myself during the course of development. Security, also – being a large part of the prototype requirements, my focus has been to a large extent on ensuring the capability grant and delegation process is robust. Usability and UX felt challenging as evaluands, however, since I could by definition not do that alone and, as already mentioned, it was a domain I was very anxious about entering.

I determined that I would ask a small number of friends to just try the prototype app and feed back any major issues. This process is dependent on the third-party ‘launcher’ app, which I had no control over. I verified that I could load and use the prototype app within ‘launcher’ on my own computer, and prepared some basic instructions for the test users I would engage. This was implemented as an official Release¹⁸, and the app file and instructions can be found on that release page at <https://gitlab.com/alexpoor/radhis/-/releases/v0.0.1>.

This immediately threw up an issue with one user who installed ‘launcher’ but experienced an issue with opening the prototype app file. This was acknowledged as an issue by the ‘launcher’ development team and, as at April 2022, is still an open issue¹⁹. Another user could not install the ‘launcher’ app at all, and had to switch to a different operating system. Ultimately, I had three users who opened the prototype app in ‘launcher’ successfully, and were able to confirm that all the base functionality worked.

One interesting finding from this process was that very long entries could be saved, which would then severely impact performance for that user – effectively crashing the app. This led to the creation and implementation of another requirement to enforce a maximum character limit on any freetext field. Holochain documentation actually states that the maximum entry size is 16mb²⁰, which would represent a very large amount of text. I did nevertheless implement a limit of 100 characters on the Content field, and 50 characters on the Resource Type field, and this problem did not reoccur.

Another finding was around the formatting of dates. The user interface requires dates to be sent to the Holochain backend in UNIX time format. UNIX time is displayed as an integer, showing the number of seconds that have elapsed since 1 January 1970. The UTC time Wed Apr 20 03:13:09 2022, for example, would be represented simply as 1650424389 in UNIX time. This format is ubiquitous in computing, and I had become accustomed to testing the prototype app by using these integer values for dates. All test users fed back that this will make no sense to users, and a human-readable date format should be used (which, in retrospect, was rather obvious). I implemented a change to fix this issue, whereby a date picker element was added

¹⁸A GitLab feature which allows you to collect files and collateral together to constitute a working package for people to use (<https://docs.gitlab.com/ee/user/project/releases/>).

¹⁹<https://github.com/holochain/launcher/issues/38>.

²⁰https://docs.rs/holochain_types/0.0.1/holochain_types/prelude/constant.ENTRY_SIZE_LIMIT.html.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

to the form which most users would be familiar with. The frontend code automatically converts the selected value to UNIX time, before sending to the Holochain backend.

Having cleared up these initial issues, I wanted to conduct a session where at least four people in total were all connected simultaneously to the app and interacting with it. This would not only flush out any other issues, but would constitute a type of basic usability test.

Usability testing is another area with a large literature base and a range of formal approaches. One of the dominant approaches in usability testing is the “think aloud” method. This method requires a user to test the product, whilst thinking aloud as they do so. An evaluator is present, and they observe the user and record observations about what they are saying and doing (Lewis, 1982).

It is important to revisit, though, what our research goal is at this point lest we enter the proverbial rabbit hole. Hertzum states that “usability testing is an activity in the process of product design. Thus, the purpose of usability testing is to inform design” (2020, p.4). As I have already stated, the focus of this thesis is not about product design – it is about proving a concept around distribution, and gauging its applicability in the real world. Even at the conclusion of this thesis, the need to definitively scope product design components for real world application is a very long way off. As mentioned in section §5.4, the user interface development for the V2 prototype is being carried out with the dual goals of increased understanding and ability to disseminate findings, not because I plan to make it available to end users in the near future.

A full-blown usability testing process can therefore reasonably be considered out of scope. However, I will get some way towards it with the group testing session I have noted above. My concept for this session is to group four people, who will be geographically remote and using different devices, to role-play a scenario featuring all of the core requirements (please refer to table 6.1) around health data. There will be a very informal “think aloud” component, where any issues can be immediately fed back and logged, but I am certainly not attempting here to conduct a robust piece of usability testing. The requirements for this group testing session are summarised in table 6.4.

These requirements are about the group testing session itself. The actual prototype requirements, which all need to be fully functional via the user interface, will all be tested as part of the scenario which is indicated as the final item in the table. The scenario simply needs to ensure that all functionality is tested and works – the other requirements in table 6.4 deal with broader issues, such as RP8 (there is a working user interface) and RP9 (the app can be easily installed and run) from table 6.1.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

Table 6.4.: V2 group testing session requirements.

Requirement	Purpose
Users are geographically remote	To prove real-time connectivity of Holochain app
Users use a range of devices	To verify device/OS compatability
Users' screens are all visible	To record and verify performance
A basic roleplay scenario is developed and utilised	To ensure that all features (requirements) are tested in a network as close to real life as possible

6.2.4.1. The group testing scenario

A scenario was constructed for four users, which would tie all the functionality in the app together into a basic story. This could have been done in a number of different ways (simply recording users testing components, or a more rigorous “think aloud” process), but I also wanted to be able to record this session and make it available as part of this thesis. In this way, there is an easily accessible method of verifying that the prototype app works and meets all the different requirements I have noted. The scenario therefore has the following structure. There are four players: an individual, their partner and two health professionals. The individual is the focus, and the scenario sees them engage with the other players. The actions taken are as follows:

1. User records some health data as a proxy for patient-generated health data
2. User reviews that data
3. User gives read permission to their partner
4. The partner successfully reviews the health information, and verifies that they cannot review anyone else's health information
5. User gives write permission to both health professionals
6. Both health professionals create some data as a proxy for a consultation of some kind
7. A health professional verifies that they cannot create data for anyone else
8. The user revokes permissions

6. Research Phases 2-3: V2 Prototype design, build and evaluation

9. Other users verify that they can no longer carry out actions they were previously able to.

Successfully achieving all the above in the group test setting will confirm achievement of all requirements in table 6.1, as well as those we just discussed in table 6.4.

6.2.4.2. Running the group testing session

The session was scheduled for 12 February 2022, and all but one participant had successfully tested the V2 prototype app beforehand. All participants were remotely located – three in Hamilton, Aotearoa New Zealand, and one in Toronto, Canada.

The session was to be conducted, and recorded, via Zoom²¹. Zoom permits all participants to share their screens on demand, and this is a feature that would work well with the scenario format described above. The scenario was conducted and recorded without any major technical issues. One participant was experiencing a large amount of fan noise from their computer, which rendered them inaudible in the video recording. This was rectified subsequently by re-recording only the audio and splicing it into the video using an editing tool²². There were some other minor issues with the video recording. For example, I had inadvertently shared the wrong screen at different points – this was also rectified by blocking out parts of the video using the kdenlive software. Subtitles were also manually added as an accessibility feature.

The final video for this group testing session is available online at <https://vimeo.com/680251596#t=949s>. This video will show a full working demonstration of remote users all connecting in a distributed network to create and share data. Users were using different operating systems (one used Ubuntu, one used Arch Linux and two used Windows 10). The app was installed independently by each user, following the instructions on the Gitlab release page.

6.3. Summary

At this stage I have achieved the following:

- Built and evaluated V1 prototype against original set of requirements (table 5.1)
- Built V2 prototype featuring a full user interface and ability for users to install and use Holochain app
- Tested V2 prototype with individuals and as a group.

²¹<https://zoom.us/>.

²²<https://kdenlive.org/en/>.

6. Research Phases 2-3: V2 Prototype design, build and evaluation

I now have a fully functional app which anyone can install and utilise, as well as a video available showing it being used as part of a roleplay scenario. In the next chapter I will explore how such a tool could be implemented within the Aotearoa New Zealand health sector.

7. Research Phase 4: Interviews

The full research plan is visualised in figure 3.1. As already discussed this broadly follows the DSR methodology, the study requirements of which are shown in table 3.4. The first of these study requirements is noted as “Understanding of user attitudes and requirements, and SME perspectives” and this encompasses: Research Phase 1, where we clarify this is a real problem that requires an innovative solution, and; this Research Phase 4, where we are concerned with locating the prototype and its implications into a real world context. That is, the working prototype does not exist in a vacuum but must be located within the appropriate regulatory, legislative, cultural and political context. The purpose of this research phase is to pose and resolve the question:

You have demonstrated that this is a problem needing to be solved, and shown a way that it could be solved. What are the obstacles to further implementation of the solution?

In this respect, the interviews pose a thought experiment and ask participants to consider the implications from their own perspectives as SMEs in fields related to health information or information in general. There is no requirement, therefore, that interview participants have already seen or used the prototype – I am seeking their thoughts on how distribution *might* be applied.

This Chapter will detail the approach taken for this research phase, as well as an overview of the findings, before everything is brought together for discussion in chapter 8. Please note that, to minimise confusion for participants, I elected to use a simplified architectural taxonomy composed of centralisation vs decentralisation. This is at odds with the discussion in section §2.3. However, it was clear to me that anything other than centralisation was a novel concept for participants and, therefore, balancing discussion with its apparent opposite (which is actually ‘distribution’) would aid understanding. Please bear this in mind when considering participant quotes in section §7.7.

7.1. Choice of method

Interviews offer the researcher the ability to gather rich qualitative data, which elucidate the participant’s experiences and worldview (Rubin & Rubin, 2012). The purpose of this research

7. Research Phase 4: Interviews

phase is to understand more about how supportive the sector is currently of implementing distribution. There are many strands to this; the dominance of centralised data models has influenced the way we see and think about things when dealing with data. Its technological immaturity has meant that it is not thought of as a viable way to build new solutions. Furthermore, there are many practices – supported by regulation and legislation – that reflect this dominance, in sometimes subtle ways¹. Of course, many of these rules are around important principles such as claiming and funding, or public health. It is likely that, if we were to ever see any move towards distribution, these tenets would remain essentially untouched. But, this is exactly the kind of discussion that we need to stimulate in this research phase with the contributing perspectives and experiences of a wide range of experts.

I will also attempt to identify the extent to which different participants see the current situation as a problem that needs to be solved, since we have already seen in chapter 4 that survey findings indicate that the aims of a RDHIS are important to them. There are many variables involved here.

Patton (2015) has noted that qualitative methods have a strength in uncovering the nuances of any local implementation, or adaptation. Rich qualitative data may be more complex to gather and understand, however it will offer much fuller insight into the ‘softer’ issues which are crucial to understand both at national and sub-national level in Aotearoa New Zealand. The choice of this method is well summarised by Patton’s claim that the point of interviews is “to elicit relevant answers that are meaningful and useful in understanding the interviewee’s perspective” (Patton, 2015, p. 471).

This qualitative data will be obtained via in-depth semi-structured interviews, and my specific rationale for this approach is as follows:

- To capture data which is currently unknown to the researcher, or otherwise not available in the literature. There is evidence to show that knowledge management and information dissemination is a particular issue in health (Bordoloi & Islam, 2012), and therefore in-depth interviews are an effective way of capturing this tacit knowledge.
- To seek expert opinion from SMEs in a range of different areas relevant to this research. For example, policy, legislation, medicine, IT and Te Ao Māori. It will not be practical to focus on any of these areas in depth as part of the primary research in this thesis, so interviewing experts in these areas is an effective way of rapidly understanding key issues.

¹For example the New Zealand Health Act (1956) stipulates a number of requirements around personal information, which all assumes that data is held centrally by a Crown entity. This works well currently, but is clearly an example of legislation that may need amending if individuals are to own and control their health data.

7.2. Participants

7.2.1. Participant selection criteria

The aim of this research phase is to understand current and future issues relating to the implementation of a RDHIS in Aotearoa New Zealand. While the focus is health information, the principle applies equally well to any other area currently served by centralised models². That is to say, the participants need not be directly involved in the health ecosystem. The health ecosystem is already complex enough, however, with a range of stakeholders from different professional groups all operating at different tiers (from primary health to community NGOs to specialist tertiary centres, and all the different specialties and sub-specialties in each). It is also important the participants, in aggregate, cover a broad range of stakeholder perspectives. This will not be exhaustive, of course, and this should be noted as a primary caveat on the findings from this research phase.

As a starting point, I have utilised a pre-existing taxonomy of health data stakeholders offered by The European Institute for Innovation through Health Data (2021) and this is shown as table 7.1.

This table shows that the taxonomy does not necessarily translate very well into the Aotearoa New Zealand context. Additionally, it misses some nuance around an individual's skill set and worldview that will be important to know for the interview process. Using table 7.1 as a starting point, I have identified the following key factors that should ideally be encompassed in the participant group.

7.2.1.1. Health sector role

There are four key roles that I want to see represented in the participant group:

1. **Individual.** This is the voice of the individual, or patient, and is critical since the aim of this thesis is to demonstrate how nascent technologies can empower individuals, and contribute to better health outcomes.
2. **Clinician.** There are many different clinical roles that are required to interact with patients and patient data. It is important to capture their view of distributing health information, based on practical experience with incumbent systems and processes.

²Not to be facetious, but technically this means *all of them*. However, there are obviously some domains of information which are more amenable to the principles of a RDHIS than others. For example, it is perhaps reasonable that the criminal justice system, or the welfare system, retain centralised approaches. Research around the possibilities of distributing these other, perhaps more controversial, data domains would be extremely useful and interesting.

³Predominantly via the Health Research Council, but other bodies exist.

7. Research Phase 4: Interviews

Table 7.1.: Health data stakeholders. Adapted from The European Institute for Innovation through Health Data (2021).

Stakeholder	Description
Patients	Individual patients or members of society
Health policy makers and funders	In the Aotearoa New Zealand context, the Ministry of Health
Standards organisations	Currently a range of standards are being used at different places in the health system
Pharma and life sciences	Organisations conducting research to make new healthcare products
Health ICT	The vendor ecosystem who build health IT products
Data driven industry	A distinct group of people or organisations utilising health data in different ways
Healthcare providers	Organisations delivering healthcare to people
Academia and research bodies	Health research is a focus of many different academic specialties and, in Aotearoa New Zealand, there is a robust system for funding and delivering research ³
Health and data strategy	Other organisations utilising health data to inform strategy

7. Research Phase 4: Interviews

3. **Technician.** This role is meant to capture anyone involved in the IT or underlying data component of healthcare. This may be people working currently with centralised systems or in a data-facing role.
4. **Policy.** This role encompasses anyone involved in research, designing and setting policy around data, especially health data.

Obviously, these roles will overlap – a Clinician will always also be an Individual, and Technicians may also be Clinicians. I should also reiterate that, within the constraints of this thesis, it is not possible to accurately represent each of these large and diverse groups.

7.2.1.2. Health sector tier

There are several tiers in health systems, each of which having different pressures, issues, systems and practices. It will be important to capture qualitative data from as broad a range as possible. I have formulated the health sector tiers as follows:

1. **Community.** This tier encompasses any community, NGO or iwi provider. These groups play a vital part in the Aotearoa New Zealand health system ranging from counselling, to housing, to mental health and addictions and are not formally part of the state apparatus⁴.
2. **Primary.** The Primary tier is the formal Primary care sector, which is mainly comprised of general practice in Aotearoa New Zealand. Ordinarily this would include the Community group but, having spent time working with data in both of these tiers, I feel strongly that there are differences in culture, service provision and philosophy that are sufficiently different to justify separation.
3. **Secondary/Tertiary.** This tier in Aotearoa New Zealand is essentially any service provided via a District Health Board and/or in a hospital setting.
4. **Central government.** Finally, the government via the Ministry of Health, are a key tier in the health system. They determine the policy settings and funding mechanisms by which the health system operates.

The tiers will only apply to participants who are working in the health system, and obviously do not apply to individuals or patients not part of any other group.

⁴These providers are typically contracted to central government to provide specific services.

7. Research Phase 4: Interviews

7.2.1.3. Cultural lens

The final group of factors, and of vital importance in Aotearoa New Zealand, is to understand the perspective of different cultures – specifically, to draw out the relevance of a RDHIS for Te Ao Māori but, if possible, also Pasifika. There are obviously other cultural groups within Aotearoa New Zealand, but Māori and Pasifika specifically have tended to have worse health outcomes and research indicates that current service delivery models have not met their cultural needs (Health and Disability System Review, 2020; Health Quality & Safety Commission, 2019)⁵. It is important therefore that this research phase should ensure there is a voice for Māori and Pasifika stakeholders.

7.2.1.4. Use of the selection criteria

The selection criteria in the preceding sections are simply meant to ensure that a reasonable cross-section of skills, knowledge and experience is represented in the interview participants. It is not a ‘tick box’ exercise and, similarly, it is important to note that this research phase is not intended to be an exhaustive study of all conceivable perspectives. It is important that there is a reasonable cross-section, and that each participant could speak with authority from their own perspective. This leads us on to a factor which is not clarified in the above criteria, and that is: the participant must have authority, or *mana*, in their field. The conclusions from this research phase are intended to illuminate the changes that have to be made to accommodate distribution of data in the health sector (and potentially beyond). Therefore, it is important that participants have sufficient experience and knowledge to talk confidently about this. It implies individuals who are in senior positions, or leading thought in their particular field.

The approach used is *purposive sampling*, which is a flexible design that utilises the researchers’ judgement (Robson & McCartan, 2016). This approach works particularly well when the researcher is involved in the field of research, and has networks or connections in it. Identifying such participants is a rather subjective matter. It is possible to be over-confident in one’s ability to identify the right people and, naturally, some good opportunities will be missed. However, I have been involved in this field for some time and feel confident that I already know the key people who can speak about the Aotearoa New Zealand context. Additionally, the networks of my supervision team also offer the ability to broaden the search and find other suitable participants.

I also leave open the possibility to utilise *snowball sampling*. This is where “participants are used to identify other members of the population, who are themselves used as informants”

⁵In fact HQSC found that this is not a question simply of *access*; “our results suggest engaging with the health system *can increase advantages for non-Māori and disadvantages for Māori* [emphasis added] across the life course” (Health Quality & Safety Commission, 2019, p. 10).

7. Research Phase 4: Interviews

(Robson & McCartan, 2016) and, indeed, this did occur during the research process.

7.2.2. Sample size

In the previous section I noted that all participants need to be suitably experienced and qualified to talk in-depth on this topic; they should be subject matter experts within their field. Because of this, the interviews were likely to take at least one hour and, in some cases, much longer⁶. Because of the other research phases presented in this thesis, a nominal guide of twelve participants was selected as being suitable.

The literature is understandably non-committal on the question of how many participants are ideal, since it is contingent on the specific thing being studied and the environment and circumstances surrounding it. Patton's assertion that "there are no rules for sample size in qualitative inquiry" (2015, p. 184) is unfortunately true, with other writers bemoaning a focus instead on instinct or subjective experience (Malterud, Siersma, & Guassora, 2016; Mason, 2010) as a means of arriving at a responsible analysis. A helpful addition to this debate is provided by Malterud et al. (2016), however, who suggest that attention should shift from the *number of participants* to the *range and depth of new knowledge* that the interview process is generating. This is referred to as 'Information Power'. Making an explicit distinction to the concept of 'saturation' (which is specific to Grounded Theory, but is utilised imprecisely across a range of other methodologies), their suggestion is simply that, where more information is held and generated, a lower number of participants is generally needed. It is useful to identify an approximate sample size as a starting point (and this can certainly be informed by experience), but very important that it is continuously reevaluated and assessment is made about the potential diminishing returns of interviewing more participants.

7.2.3. Participant engagement

Following this approach, I made an initial shortlist of 12 individuals who I had a connection with, and would meet the participant criteria discussed in section 7.2.1. This was discussed with the supervision team, and additional candidates were identified from their networks. This stage utilised the principle of *purposeful sampling*, whereby my personal knowledge of the individuals (or the knowledge of my supervisors) could guarantee that they had pertinent knowledge and experience and the interview would generate rich data. Patton notes that the success of this approach, from a research integrity perspective, relies on "selecting information-rich cases for study in depth" (2015, p. 169). This aligns well with the concept of 'Information Power', where

⁶In the final analysis, three interviews took more than one hour with a maximum interview time of one hour and forty five minutes.

7. Research Phase 4: Interviews

it is contended that interview design should focus on the value and depth of data or information gathered.

With an initial list of 20 candidates, I proceeded to make contact and seek their *in principle* consent to take part in the research; 15 responded positively. After AUT Ethics Committee approval was given on 17 May 2021 (please see section §7.4 for further information), I followed up with these candidates and was able to re-engage with only 11 – a phenomenon that is well understood in the literature (Robson & McCartan, 2016). Two additional participants were identified via snowball sampling, bringing the number of consented participants up to 13. A summary of interview participants is shown in table 7.2, where participants are identified (with their consent) for the purpose of establishing expertise and relevance of findings.

Whilst I could be very confident that interviewing these participants would generate rich and useful data, it is important to check how well the mix of participants fulfills the selection criteria (section 7.2.1).

7.2.3.1. Assessment of participants against selection criteria

The selection criteria encompass three domains: health sector role, health sector tier and cultural lens. In figure 7.1, a visual representation is shown against the health sector role and tier domains only. This information was gathered as part of the interview process, as well as my own knowledge or publicly available information about each participant. It is certainly not intended to summarise this group of peoples entire lived experience into a simple table; it is merely a heuristic to gauge where any potential ‘blind spots’ may be present in the research outputs.

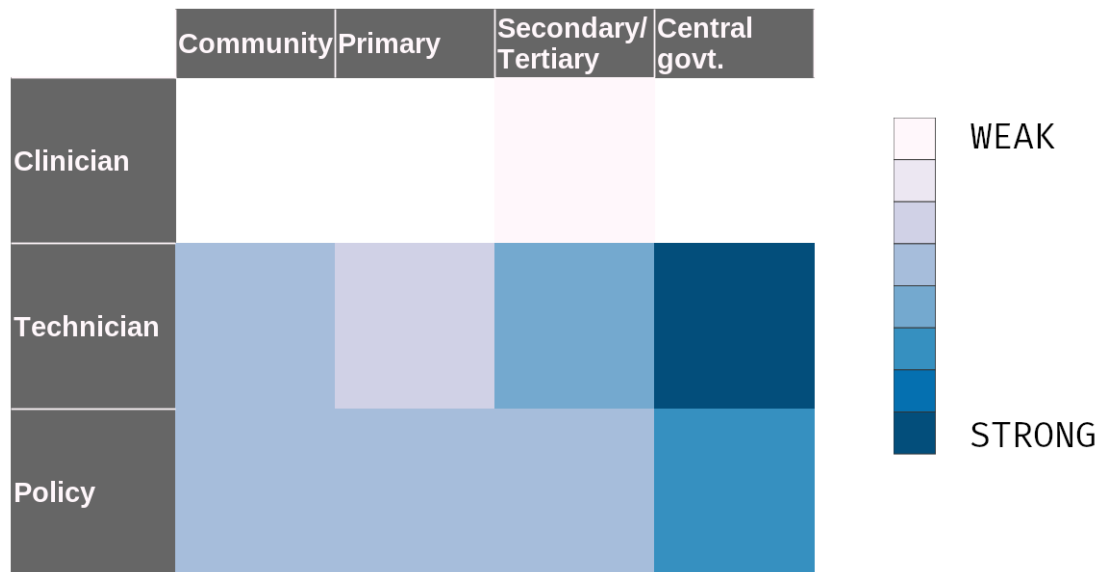
7. Research Phase 4: Interviews

Table 7.2.: Initial participant list and selection criteria

Name	Role
Dorothy Adams	Seconded to the OECD. Formerly Chief Executive, Social Wellbeing Agency; GM Data, Evidence and Insights at the Ministry of Social Development
Dr Andrew Chen	Research Fellow at Koi Tū – the Centre for Informed Futures at the University of Auckland
Mark Corbitt	Chief Executive, Patients First
Wendy Hamilton	Chief Data Steward, Ministry of Education
Shayne Hunter	Deputy Director-General Digital and Data, Ministry of Health
Leaupepe Rachel Karalus	Chief Executive, K’aute Pasifika Trust
Joy Liddicoat	President – InternetNZ, Human Rights and AI researcher (University of Otago) and formerly Assistant Commissioner, Office of the Privacy Commissioner
James Mansell	Senior leader and Special advisor encompassing multiple government portfolios, focusing on use of data to improve effectiveness
Dr Will Reedy	Chief Executive, Spark Health, and medical doctor
Kylie Reiri	Chief Executive, Nicholson Consulting
Ngapera Riley	Chief Executive, Figure.NZ
Simon Ross	Chief Data Steward, Ministry of Health
Steve Creed	Former senior leader across a range of health and public service settings

7. Research Phase 4: Interviews

Figure 7.1.: Assessment of participants against selection criteria



The visualisation shows that there is good coverage across the Technician and Policy roles, but poor coverage across the Clinician role. The initial participant list had identified two key people operating at a senior level clinically, both of whom unfortunately were in the group of participants that I could not re-engage with following AUT Ethics Committee approval. Given the schedule of this thesis, it was not possible to identify and engage with a suitable alternative within the available time. The clinician voice is therefore noted as being under-represented in this research phase. This is certainly not ideal, but it should be remembered that this phase is not intended to be a canonical representation of current thinking about distribution. It is the first step in a journey towards identifying how distribution *might* be possible in the Aotearoa New Zealand health sector and there are, nevertheless, important insights from those who did participate. The under-representation of the clinician voice is noted in section §8.3 as an area requiring further research.

What figure 7.1 also shows is a slight imbalance towards both the central government and hospital tiers of the health sector. This is not necessarily a bad thing, since these are the tiers with the most influence and power within the health sector, as it stands at the time of writing. The Ardern government announced a wide-ranging reform of the NZ health system in 2021, with the tacit acknowledgement that the Community and Primary tiers have not received enough attention. This also signalled the dissolution of the District Health Board structure, which has managed the flow of funds from its Vote Health allocation down to Community and Primary providers (with arguably a disproportionate focus on hospital services).

7. Research Phase 4: Interviews

I removed the ‘Individual’ role from the assessment shown in figure 7.1, since every interviewee also effectively speaks as an individual – and many participants did indeed speak from the perspective of themselves as health system users. Again, it is unlikely that they would be able to represent the full range of individual experiences of the health system – particularly, by definition, those who are marginalised from it – but this has been touched on in Research Phase One (please see chapter 4), and is an area that will benefit from further research.

7.2.3.2. Cultural lens

The cultural lens criterion does not fit neatly into a visual representation, and is also something that would be insensitive to quantify. Individuals have complex relationships with their culture and ethnicity and, as a researcher, it felt improper or intrusive to ask participants directly about it. Participants did speak willingly about this, but I wanted their inclusion to be solely about their *mana* within their field.

One participant, who has been active within Te Mana Raraunga (the Māori Data Sovereignty Network), highlighted the complexity of this issue by specifically requesting that their participation is on an *individual* basis and that their comments should not be misrepresented as if they were speaking on behalf of the broader group. Firstly it is somewhat ironic – and betrays my own cultural bias towards individualism – because I had never intended to do so. But, secondly, it speaks to perhaps a type of *engagement fatigue* (The Office for Māori Crown Relations - Te Arawhiti, 2018) whereby indigenous people and ethnic minorities are invited to speak on behalf of their entire heterogeneous and multifaceted cultural group – or, worse, as a sole representative simply in order to tick a particular box.

In summary, there certainly has been a cultural lens applied to this research phase via the contribution of some participants. This is not to say that full engagement or consultation has been done with Māori – it has not, and that was never the intention. But some important cultural issues certainly arose from these interviews, and they will be discussed in section §7.7.

7.3. Interviews: Design and implementation

Having clarified the process around participants, we now turn our attention to the design and implementation of the interviews themselves. As with all research phases, the interviews must directly link to the *purpose* of this research, and the research questions themselves. That is, the interview process must be wholly congruent with the overall research aims (Jones, Torres, & Arminio, 2013). This section will outline the design and implementation processes for this research phase.

7. Research Phase 4: Interviews

7.3.1. Developing the interview protocol

In designing the overall protocol for this research phase, I will rely on the Interview Protocol Refinement (IPR) framework developed by Castillo-Montoya (2016). This framework proposes the completion of four key phases, before being ready to interview participants:

- Alignment with research questions
- Constructing an inquiry-based conversation
- Receiving feedback on protocols
- Piloting the interview protocol (Castillo-Montoya, 2016).

These phases will be discussed in the following subsections.

7.3.1.1. Checking alignment with research questions

The interview process must be carefully designed, to ensure that it is congruent with the research aims and therefore leads to the collection of useful data. Without this step, the interview may be nothing more than a pleasant conversation.

The research questions are presented in section §1.3 and it is clarified there that, although there are three questions for the thesis to answer, they are encompassed by three separate research phases. Therefore, while Castillo-Montoya (2016) asserts that the interview process must relate directly to all research questions in order to be useful, the structure of this thesis as a whole means that this need not be the case. For example, some research questions can be answered via the literature review, or via prototype development. Only one research question is wholly the responsibility of this interview research phase: *How could a distributed health information system be implemented in Aotearoa New Zealand?*

To answer this question we necessarily have to consider the current state and then think about those things which represent meaningful obstacles. Some of these may be technical (the capability of extant IT systems in the health sector), some may be political (anxiety about losing access to foundational data sources) and yet others may be legal (for example, Health Act 1956 stipulations about recording and managing patient information). Another way to look at the question may be to ask instead – *how far has centralisation influenced structures and processes that sit around the health sector, and health data, and tend to further embed it?*

The interview questions should therefore try to elicit these obstacles, as well as to unravel the centralisation problem, so that the information can be analysed in aggregate and presented for further discussion. I therefore decided to use the following questions for the interview guide:

7. Research Phase 4: Interviews

1. What advantages could we realise by having individuals own and control their health data?
2. Can you think of any legal or regulatory obstacles?
3. Can you think of any cultural or social issues?
4. How do you think incumbent systems and providers would respond?
5. How do you think the broader health system would respond?
6. Is there anything else you'd like to talk about, that we haven't already touched on?

The first question is designed more as a conversation starter, to start thinking about the decentralisation concept in more practical terms. Rubin and Rubin (2012) note that, while it may feel efficient to focus purely on key questions, there is necessarily a very human element to the interview process and, therefore, it is important that some rapport is first built before moving on to questions that are more directly linked with the research purpose. This first question is of course directly linked to the research purpose anyway. However, it is a very broad open question and one that most participants would find it easy to have an opinion about.

The following four questions, two to five, are clearly more focused on the elicitation of those obstacles which may have been embedded by the dominant centralised model. These questions provide opportunity to discuss how to overcome identified obstacles, and move towards a distributed model. The final question is a closing question, giving the participant an opportunity to make any additional comments or observations they wish.

7.3.1.2. Constructing an inquiry-based conversation

If it is so important to centre the interview on the research questions, then why not simply use these in the interview? The IPR framework here calls us to ensure a clear distinction between research questions and interview questions – “research questions formulate what you want to understand; your interview questions are what you ask people to gain that understanding” (Maxwell, 2012, p. 101). Creating an ‘inquiry-based conversation’ therefore requires the researcher to employ user-friendly language, avoid jargon, and to consider social norms and the contexts of the participants themselves. Interview question design is essentially a process of translating the research questions, in order to get maximum engagement, clarity and safety for the participants.

Having said this, the interview should not necessarily be a facsimile of an ordinary social conversation. Something I noted while carrying out the interviews was my urge to step in and make a point, or to correct something I felt was incorrect or misunderstood – effectively to engage in

7. Research Phase 4: Interviews

convivial *debate*. However, the interview is not about persuading people, or promoting something you believe in. The focus must be to understand what the participant thinks and feels about your subject. In this way, it requires constraints that are not found in ordinary social conversation. Castillo-Montoya (2016) notes the example of following up with “why?” questions. In an interview setting, this could be perceived as judgmental or antagonistic whereas it is perfectly acceptable in social conversation.

Striking this careful balance between inquiry and conversation can be aided by thinking explicitly about the *types* of question being asked, and when to utilise them. For example, Rubin and Rubin (2012) identify four types of interview question, each playing an important role in striking this balance:

Table 7.3.: Interview question types. Adapted from Rubin and Rubin (2012).

Question type	Description
Introductory	Neutral and non-threatening questions, eliciting general information
Transition	Link introductory questions to key questions
Key	Directly connected with research questions and study purpose
Closing	Easy to answer and move the interview towards closure

Introductory questions establish rapport, and help to move the participant’s mindset towards inquiry. Transition questions will pick up on something the participant has said, and use it to move on to another topic or a Key question. Key questions themselves are the ones I noted in the previous section, and are most directly linked to the research questions. Closing questions should provide an opportunity for the participant to discuss anything else they want and to start transitioning out of the inquiry space.

7.3.1.3. Receiving feedback on interview protocols

An evaluation step in protocol design is important, in order to make sure that collected data is as useful and robust as possible. It is also important that some testing takes place to ensure that questions are likely to be correctly understood and interpreted by participants. Hurst et al. have noted that failing to do this introduces “the risk of later collecting invalid and incomplete data”, but also to bear in mind that doing so “is not a guarantee of ... success” (2015, p. 827).

This evaluation step should incorporate consideration of a range of dimensions, for example: structure of the interview protocol; validity of questions being asked, and; congruence between interview and research questions. In practice this involves a read through the interview protocol with another person, who should ideally share some characteristics with the actual participants.

7. Research Phase 4: Interviews

This was not practical in my case, due to the narrow group from which participants were drawn – but also because they would all be speaking from such varied perspectives. Nevertheless, I did perform a basic read through with a colleague in the health sector and this led to some changes around jargon and terminology⁷. This was an important change, since I am proposing to interview participants from a wide range of technical and non-technical backgrounds. If I am speaking with a technical expert I can, of course, flex my approach to use more technical language. But the protocol, used as a guiding framework, should not embed any assumptions in this area.

7.3.1.4. Piloting the interview protocol

Following the previous three steps, the researcher would have a protocol which they could be confident is congruent with the research purpose (and research questions), and which balances the requirements of an ‘inquiry-based conversation’. Some improvements will have been made, based on testing and feedback in the prior step. This step takes that even further, and proposes a full simulation of an interview under conditions which are as realistic as possible. In my case that meant hosting an interview using the protocol, via video conference. In this stage we are more concerned with the practicalities of utilising the interview protocol, as opposed to focusing on participant understanding or thought process (Castillo-Montoya, 2016). This step enables the researcher to gain a practical sense of how long the interview would take, as well as to understand any practical or logistical issues that may occur.

I conducted pilot interviews with two colleagues who both had strong health sector, and research and evaluation, experience. However, they did not have any specific knowledge (technical or otherwise) on the research topic of distributed data (which, in fact, would be true for most interview participants). These interviews were done using a software video conference tool (Zoom), which allowed the audio and video to be recorded. With only two participants, there was also no technical limitation on the allowed time. Because the interviews were to take place in the last three months of 2021, I reasoned that most people would be very familiar with using these tools given the lockdown constraints imposed since early 2020 when the Coronavirus pandemic emerged. The pilot interviews resulted in the following findings:

1. **Interview duration.** The pilot interviews took between 35 and 40 minutes to complete. I had expected them to take around one hour, and so this surprised me. However, on reflection I considered that the pilot interview was only a simulation and, thus, it is possible that the participants picked up that the breadth and depth of their responses were not the focus.

⁷Specifically that I had been referring to the centralisation/decentralisation dichotomy, on the risky assumption that participants would have carefully digested the Participant Information Sheet.

7. Research Phase 4: Interviews

It is also true that the individuals involved do not have a strong stake in this field, and so their answers were mainly from a personal perspective (which did make some of the questions hard to pilot). Based on my personal knowledge of the interview candidates, I felt confident that the interviews could easily take one full hour and, in some cases, possibly even longer⁸.

2. **Introduction and closing.** Upon trying to provide a realistic simulation of an interview, I realised that I did not have any guide as to how to open the conversation. Because of this, the first part of the pilot interviews was quite rambling on my part, and I did not feel confident it gave a good impression to the participants about my integrity as a researcher. This is perhaps a more subjective point, but I knew that rapport was important to generate useful data (Rubin & Rubin, 2012) and rapport could be hampered if the participant feels the interviewer is not prepared, or not in control of the process. I therefore implemented a standard opening statement into the interview protocol which, although it was largely a reinterpretation of content from the Participant Information Sheet, would serve as a concise focal point before entering into discussion. To a lesser extent the same was true of the closing statement. There are some key facts that I needed to reiterate to participants (for example, next steps around transcript approval) and this should be consistent.
3. **Management of apps and notifications.** A practical learning from this process was around management of notifications on my device. Since the video conference tool ran from my laptop, where my audio and video would be shared with the interviewee, there was a possibility that this could cause some interruption. Indeed, this did occur, as some app notifications on my laptop made audible sounds which were distracting during the pilot interviews. I therefore added a protocol step around deactivating any notifications prior to an interview. These notifications are now so ubiquitous that the participants themselves will likely experience them. However, I wanted to ensure that any such distraction was not generated from myself as the interviewer.
4. **Understanding of network types/taxonomy.** The taxonomy presented in section §2.3 describes three network types – centralised, decentralised and distributed. The research focus here is on the distributed type. I found during the pilot interviews that this became very confusing for the participants, who easily understood ‘centralisation’ but struggled to distinguish between the other two types. Because it was important for me during the interviews to contrast what I am proposing against the dominant paradigm (centralisation), I decided that simply presenting the opposite of this (decentralisation) would be

⁸As already noted, this is indeed what happened, with three interviews taking longer than one hour, but most were around the one hour mark.

7. Research Phase 4: Interviews

conceptually simpler for interviewees. When I describe ‘decentralisation’ during the interview process, however, I actually describe ‘distribution’ (as can be seen in the Participant Information Sheet at appendix E). This is simply because the polarity of centralisation-decentralisation is easier to understand for non-technical participants.

7.3.2. Participant interviews

The interviews generally took place within business hours. However, one participant was based in Europe and it was not possible to arrange a convenient time during the working day of either party⁹. As already mentioned, they all took place via video conference software which allowed unlimited meeting time between two participants, and the ability to easily record the audio for transcription.

Luckily there were no unexpected curtailments of interview time, or clashing appointments. This was particularly lucky given that I was interviewing senior people who would be pressed for time. Similarly, there were no technical issues with using the video conference tool. This seems to bear out the assumption I previously noted, that participants would be very familiar with it having worked through the Coronavirus pandemic for more than a year by that point.

Three interviews had to be rescheduled ahead of time, due to pressing commitments – one participant needed to reschedule on four occasions. However, good notice was given and it was very easy to be flexible and arrange alternative times.

In total, 805 minutes of interview time was used across all 12 participants. The average interview duration was 67 minutes (range 60-105).

7.4. Ethics

AUT Ethics Committee approval was given on 17 May 2021. The protocol submitted in the application notes that participants must give fully informed consent, and also being clear that all interviews are recorded and transcribed (and the mechanics and privacy issues associated with that). The approval letter is attached as appendix D.

Interview participants, upon confirming their availability and willingness to take part, were each sent a Participant Information Sheet which explains in some detail the purpose of the research and what is expected of them during the interview. Other information is also provided around retention of information, their right to withdraw and contact details of the Primary Supervisor in case they have any issues around the research process. A copy of this is attached as appendix E.

⁹This particular interview took place at 8pm CET and 7am NZDT.

7.5. Transcription

While transcription is not necessarily part of the formal research it is, of course, absolutely vital if any value is to be derived from the interviews. Transcribing results in a documented record of the conversation – one which the interviewee has a chance to amend and which can be used for a range of textual analyses. Because of the amount of audio I had garnered (and the overall workload of this thesis), I decided to make use of a transcribing service. Whilst I had used overseas, and even automated, transcription services before, my impression is that the quality is either low or – at best – many Aotearoa New Zealand terms and acronyms will be incorrect or missing. I therefore ensured that the transcription was done within Aotearoa New Zealand, and this bore fruit when additional complexity, such as local acronyms and use of *te reo*, was accurately transcribed.

All transcripts were, however, reviewed in detail by the researcher. Although this was still time-consuming, I believe it is absolutely vital to have a very intimate knowledge of the transcripts so that the coding process can be more effective. In total the transcripts generated 70,600 words, with an average of 5,900 words per transcript.

7.6. Analysis

Qualitative data analysis, particularly of interview transcripts, relies first on a coding process so that sense can be made of the rich unstructured data. We have already discussed a basic content analysis coding process in section 4.8.2, when reviewing the survey data in research phase one. Coding involves classifying the data in the transcripts and then using it to find patterns or themes. The goals of the coding process can be described as “[enabling] you to organize and group similarly coded data into categories or ‘families’ because they share some characteristic” (Saldana, 2015, p.8). That is, transcripts containing hours of conversation with interview participants, are carefully coded to tag particularly meaningful or insightful comments. We attempt to bring structure to the unstructured data.

There are, of course, many differing methodologies and frameworks for conducting qualitative data analysis. Mackieson et al. (2019) identify two main strands:

1. *Content analysis* – organising and quantifying the data into categories.
2. *Thematic analysis* – interpreting patterns of meaning in the data.

The first can be seen as a relatively simple descriptive process of coding data and analysing results and, indeed, was employed during research phase one. The second goes much further, and invites greater consideration of themes, subthemes and the generation of a coherent story

7. Research Phase 4: Interviews

from the data. All approaches rely on this foundational coding step. At the time of writing, automated coding tools are not widely or easily available, and consumer products tend to be focused on specific tasks such as call centres or sentiment analysis¹⁰.

What this means is that the task of identifying codes, and coding the data itself, is generally the job of the researcher. This introduces the possibility of biased, subjective and non-reproducible outputs. For example, as a researcher, I can choose to deprioritise coding of points which do not support my hypothesis and, conversely, I could ‘over code’ the data to in order to back up a line of argument. Of course, this would be hugely unethical but it is nevertheless a real possibility. This effect is well acknowledged in the literature; Bierema et al. (2021) found that a panel of disciplinary experts in education research exhibited quantifiable bias when manually coding data. Even when automated coding becomes ubiquitous, we must be conscious that the machine-learning and AI models those products use would be predicated on human-coded data – the bias persists and, in fact, can become reinforced (Agarwal & Mishra, 2021).

So what can be done? Saldana (2015) suggests a cyclical coding process whereby coding is performed, evaluated and reviewed. This can be done over several cycles and thus promotes a very thorough and diligent coding process. However it obviously still cannot rule out the influence of cognitive bias alone, particularly if the same researcher is performing each cycle¹¹. Perhaps it is better to acknowledge the bias, and actively centre the research process around it?

This is what Braun and Clarke (2006) advocate in their writing on reflexive thematic analysis. That is, bias simply cannot be eliminated so we must seek instead to be conscious of it and let it contextualise the research. This is very much reflected in the authors’ distaste for the notion that themes passively ‘emerge’ from the data, or that themes are waiting to be discovered. This implies a degree of objectivity – or *empirical truth* – that simply cannot exist, and it “denies the *active* role the researcher always plays in identifying patterns/themes, selecting which are of interest, and reporting them to the readers” (Braun & Clarke, 2006, p.80). Reflexive Thematic Analysis (RTA) acknowledges these decisions and intervention points, and helps the reader understand the researcher’s interpretation of the data. A six phase process is proposed by the authors¹², and I shall rely on it for this research phase:

¹⁰For example, Microsoft’s ‘Text Analytics’ service or AWS ‘Comprehend’ (<https://aws.amazon.com/comprehend/>).

¹¹It must also be noted that bias will already have affected my choice of research questions, interview participants, research methodology and choice of breakfast. However, it seems to me that the coding process is particularly vulnerable to bias, and the amount of literature on the topic would seem to support that.

¹²Although it is helpful to remember that the qualitative research process should not be prescriptive, and will rarely be linear (Patton, 2015).

7. Research Phase 4: Interviews

Table 7.4.: Reflexive Thematic Analysis process. Adapted from Braun and Clarke (2006).

Phase	Description
1 – Understanding the data	Become as familiar as possible with all aspects of your data
2 – Initial coding	Organise your data into meaningful groups
3 – Search for themes	Group codes into overarching themes
4 – Review themes	Ensure that data within themes is coherent, and themes are identifiable and distinct
5 – Define and name themes	Identify the story of each theme, and how they relate to the research questions
6 – Discussion	Provide a coherent, logical and interesting story from the data

7.6.1. Understanding the data

Having been part of the collection of data for this research phase, I have a natural advantage here. I was present during all the interviews, and subsequently listened to all of them after completion. This gave me a good sense of the content overall. Although I did not personally transcribe the audio, I did manually review each transcript and therefore have a very good holistic understanding of the data *corpus*.

Braun and Clarke (2006) promote this step as an opportunity to inductively generate some initial ideas around themes. This helps to facilitate an iterative process between theme identification, and code generation (where a theme is broader and is usually composed of a number of codes). Whilst I certainly had my own ideas in mind – as I was carrying out the interviews, and later – the interview process did throw up some unexpected initial themes¹³. In table 7.5 I set out the themes that I noted during this phase. These are mainly just general reflections on key themes as a first pass, serving to help me focus on the content of the data before going further in depth. I have included a third column – ‘expected’ – simply to indicate whether or not it was in my mind at the outset of this research phase, and this can be read as a marker of an initial theme that I found surprising. Please note, at this stage no actual coding has been performed. This is simply an informal summary of things that I noticed initially from the process of familiarising myself with the data.

¹³And, arguably, this research phase would have been unsuccessful and suspicious if that had not occurred.

7. Research Phase 4: Interviews

Table 7.5.: Initial themes from interview transcripts

Theme	Description	Expected
Risk/liability	Decentralisation limits the volume of data for central agencies to manage, and therefore reduces risk and liability for adverse events.	No ¹⁴
Integrity	There are serious issues currently with the amount of duplicate data in the sector, and the need to reconcile that and validate integrity of it.	Yes
Aotearoa New Zealand health reforms	The health structure reforms – including the creation of Health New Zealand (HNZ) and Māori Health Authority entities – raise questions about the future health data ecosystem in Aotearoa New Zealand. Specifically, what will be the data management/sharing interface between HNZ and the Ministry of Health?	No
User control	The ability for people to decide who has access to their data.	Yes
Trust/social licence	Sentiment around use of personal data, and what sort of mandate agencies feel they have around data practice and initiatives. Also concerns about trust between government funding/commissioning agencies and providers ¹⁵ .	Yes
Access to data	Specifically the concern that decentralisation will result in a reduced ability for government to access people's data.	Yes
Geopolitical data sovereignty	A focus on where the data physically resides, rather than an ability to exert control over it.	Yes
Ownership	The difficult concept of claiming <i>ownership</i> over data – particularly in a <i>te ao Māori</i> context.	No
Capability	Concern about the capability of stakeholder and individuals to operate in a decentralised system.	No
Units of aggregation	Some groups may prefer to exert sovereignty over data at different levels – e.g., as a community, <i>iwi</i> or <i>hapū</i> , rather than as an individual.	Yes
Novel data	Given social licence, control and consent – agencies potentially can access new and more meaningful data under decentralisation.	Yes
Legislative conflict	A claim that the overarching Privacy Act 2020 is enabling around data sharing – and would be amenable to decentralisation – but sector-specific legislation (e.g. Education Act, Health Act) is not.	No

7. Research Phase 4: Interviews

What this demonstrates is an indication about my own biases around the themes I would find, and makes it clear which ones I had not previously considered. This is a very helpful reflexive practice to undertake prior to formal coding since it surfaces those preconceptions and ensures that the big picture, of what *I* see in the data, is not overlooked.

7.6.2. Initial coding

The next phase entails the practical work of assigning codes to data. With the initial themes in mind from the previous phase, it is easier to understand how each item may form the basis of a theme in subsequent phases. Having said that, it is important not to let the initial themes constrain the coding. Braun and Clarke recommend to “code for as many potential themes/patterns as possible (time permitting) – you never know what might be interesting later” (2006, p.89).

This process resulted in 418 data *extracts* which were coded from the wider corpus, using the open source software QualCoder¹⁶. Extracts may be thought of as a feature of the data which appears interesting, and can be analysed in relation to the phenomenon being studied. In practical terms, they are interview excerpts that raise pertinent or interesting points about the area of study. The application of codes categorises these excerpts, which can then be *interpreted* in a later stage. A total of 85 codes were applied across the 418 excerpts¹⁷. In truth, I was surprised at how many were generated and, as the list of codes grew, it became increasingly difficult to recall if I had used one before.

However RTA does require coding at a much finer grain of detail, compared with the simpler content analysis approach utilised in section 4.8.2, so this could have been foreseen. The rationale behind this step is to ensure that we do not jump straight to higher level themes, and thereby potentially miss something interesting in the data.

In order to bring to life the ‘reflexive’ component of RTA, I will comment here on some of the issues I noted during this initial coding process.

7.6.2.1. Defining codes

This represents an issue with the level of detail at which RTA encourages you to operate within. For example several comments related to different aspects of ‘control’ over data which, in a simpler content analysis, may have been captured under a broader category of *control*. However the RTA process demands that this is split into more detailed and specific codes. For example,

¹⁴To be clear, this is one of the indirect benefits of decentralisation I was aware of, but I certainly didn’t expect an interview participants to suggest it.

¹⁵On the latter point, the Office of the Privacy Commissioner’s ‘Inquiry into the Ministry of Social Development’s collection of individual client-level data from NGOs’ (Privacy Commission, 2017) is especially pertinent.

¹⁶Retrieved from <https://github.com/ccbogel/QualCoder/releases/tag/2.8>.

¹⁷The full list is available in table F.1.

7. Research Phase 4: Interviews

the code “control over data improves health” was created based on a comment that control can make people more engaged and allow them to manage their own health more effectively.

A separate comment – that people should generally be able to have more control over their data – was coded as “people should have control over their data”. Finally, another comment stated that affording people control over their data gives them more *agency* generally. This was separated out into yet another code, because it had a slightly different focus in term of saying that it resulted in agency, and this is a concept that overlaps more clearly with Māori data sovereignty principles.

Another good example of the complexity and subjectivity in coding is an interviewee relating an anecdote about senior Māori leaders affirming that data is a *taonga*. However they also recognised that government needed to access data to properly discharge its duties. While this comment could be coded under something general like “government need access to data”, it did introduce a new dimension about the tension between government respecting Māori data sovereignty and simply needing access to data.

7.6.2.2. Internal consistency of transcripts

An additional, somewhat unexpected, complexity of the interview process is that participants didn’t necessarily present a consistent perspective on the topics and sub-topics discussed. Braun and Clarke note this phenomenon explicitly, and warn that “no data set is without contradiction, and ... [one] ... does not have to smooth out or ignore the tensions and inconsistencies within and across data items” (2006, p.89).

A basic example of this is where one interviewee expressed concern about whether individuals would have the capability to manage a decentralised system¹⁸, but later concluded that people are already having to deal with digital complexity already and decentralisation will probably not be very different – “*It’s incredibly complicated already, which makes people throw up their hands and give up. Anything you do will be better than that.*”

7.6.2.3. Evolving understanding of participants

Similarly, I noticed some participants change perspective during our interview almost as if our discussion had helped them work through the relative advantages and disadvantages as part of the process itself. One interviewee, for example, began the interview asserting that centralised data had to be maintained but, towards the end, was later commenting that decentralisation was the most equitable approach overall. This is naturally hard to quantify, but I certainly did *feel*

¹⁸Also note that we had not discussed any specific mechanics involved in a decentralised system – many of these concerns seem to have been assumed by participants.

7. Research Phase 4: Interviews

throughout the interview process that some participants came to understand the possibilities for decentralisation/distribution, simply by virtue of having an opportunity to think and talk about it.

This phenomenon is well known in computer science as the ‘Rubber Duck’ parable, described by Hunt and Thomas (2000). In it, a programmer is stuck on a problem and cannot proceed. By holding a rubber duck¹⁹, and explaining to it what is happening, the programmer can invariably arrive at the solution themselves. So there is a process where the verbalisation of an issue encourages structured thinking, which can then lead to understanding and a potential solution (if a problem exists).

I cannot locate any other literature on this topic which discusses interview participants’ deepening understanding of a topic – or shifting their opinion – by virtue of discussing it. This may very well be because interview participants are generally selected for their expertise and, presumably, this scenario would not ordinarily arise. All of my interview participants certainly were experts in their own domains but, for many, decentralisation itself was a new concept.

In fact, in places, this general understanding of decentralisation was problematic and led me to define a code “Conflating access to data with sovereignty”. This was interesting because it was not strictly a product of what the participants said. It was *my interpretation* of what they were saying, that led me to believe that some key concepts were being conflated. In this case, there were repeated statements such as “decentralisation would be good if it meant greater access to data”. These concepts are certainly related, but decentralisation is quite distinct to ‘access to data’ which could easily be achieved in the extant centralised paradigm.

If my feeling has any basis in truth, it is probable that this is because of the general lack of understanding around decentralisation/distribution I have discussed previously, and which is neatly summed up by another interviewee thus: “*I think we don’t know the full extent of what the opportunities [of decentralisation] are yet, because we don’t know that this world exists*”.

7.6.3. Search for themes

The long list of codes discussed in the previous section must now be grouped into higher order themes. Once again, there is a great deal of subjectivity involved in this phase but RTA helps to maintain transparency around the decision-making process. Braun and Clarke note that the focus of this phase is “to analyse your codes and consider how different codes may combine to form an overarching theme” (2006, p.89). However we must not only think about the relationship between codes, but between themes and any subthemes they contain. Some may naturally become themes of their own, some may build a subtheme and it might be determined that some

¹⁹The book does not explain the significance of the rubber duck specifically; it could just as well be any other inanimate object.

7. Research Phase 4: Interviews

are no longer relevant. But this phase is predominantly about grouping codes into themes, and we should be careful not to rush into organising a finalised set of themes just yet.

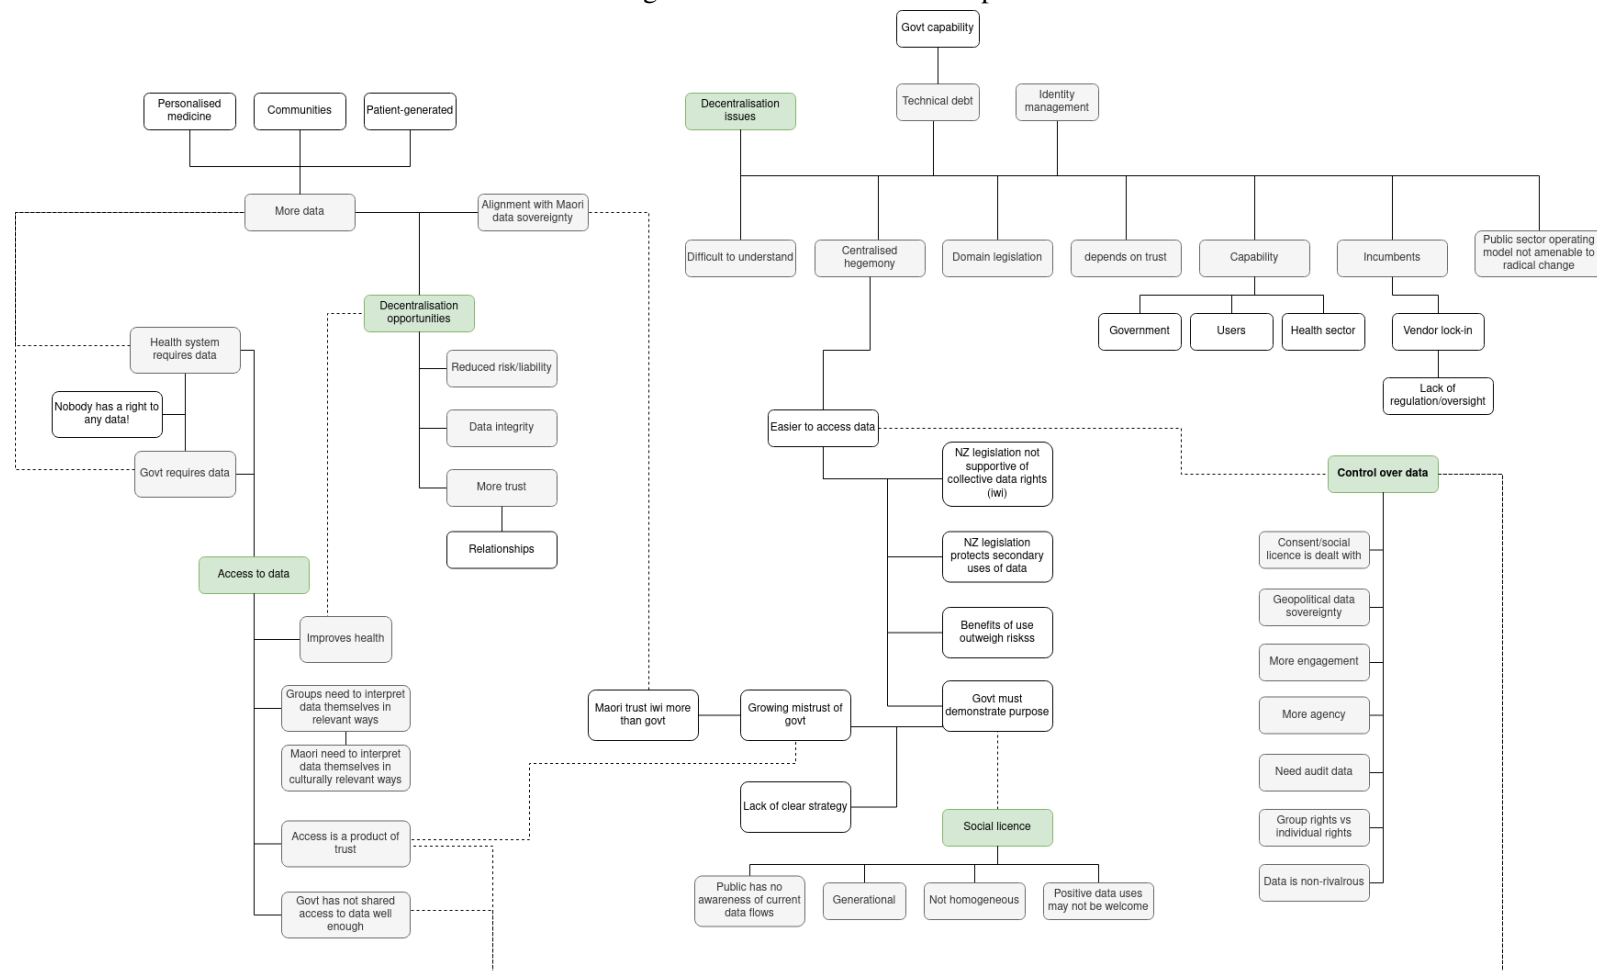
The first pass at mapping themes and subthemes from the initial coding is shown in figure 7.2. It can be seen that the map is rather large and busy, and probably is not very accessible in this thesis format²⁰. What it shows, though, is the identification of five main themes (these are shaded green in the image), and 32 subthemes (shaded in grey).

A small number of codes were discarded, either because I could justify encompassing them under a subtheme or because I deemed them irrelevant. For example, I coded nine excerpts in the previous phase on the topic of consent. The sentiment of these excerpts was that consent is being done poorly at present, and we need to think about improving its granularity and dynamism²¹. This is a very valid point in the context of the *current state*, however it is a problem that is completely solved by decentralisation/distribution (at least, as it is envisaged in this thesis). Any such ecosystem will have built in user control as a key feature; dynamic and granular consent would be a solved problem in this paradigm.

²⁰A full-resolution version of the image can be accessed at <https://imgur.com/a/9KEy9Jr>.

²¹As opposed to the current state of signing a consent once, and it never being reviewed or updated again.

Figure 7.2.: Initial thematic map



7.6.4. Review themes

This phase begins a process of refining the themes and subthemes identified in the previous section. This is essentially an *evaluation* phase, which demands reflection on the initial thematic map (figure 7.2) and testing of whether the themes have internal consistency, and whether the entire thematic map is representative of the meaning found in the *corpus*.

Patton (2015) provides a useful rubric for assessing validity – do the themes have *internal homogeneity*, and *external heterogeneity*? That is, the subthemes and codes within a theme must all directly relate to that theme. Similarly, we should ensure that the themes are distinct – across the *corpus* they should be heterogeneous, and not replicate or overlap.

Starting at the level of theme, I was reasonably confident after the initial coding process that the main themes were indeed distinct. In fact, I was even concerned that I had been too aggressive in finishing with five themes when the first phase of this RTA process (table 7.5) found at least twelve themes, and I had informally noted many others during the initial coding phase.

Nevertheless, the five themes shown in figure 7.2 are certainly distinct:

1. **Access to data.** Describes issues or possibilities around accessing data. Note that, even though I had raised focus on this as an issue when it is conflated with data sovereignty or decentralisation (discussed in section 7.6.2.3), it remains a valid issue in its own right.
2. **Control over data.** Encompasses codes dealing with a perceived right for groups to control data, or the impacts of being able to do so. Quite distinct from simply being able to access data, this is specifically about exerting sovereignty over one's data.
3. **Decentralisation issues.** Perceived issues with implementing decentralisation.
4. **Decentralisation opportunities.** Potential benefits and opportunities of implementing decentralisation.
5. **Social licence.** A theme which describes the “permission ... [agencies have] ... to make decisions about the management and use of the public's data” (Statistics New Zealand, 2018, p.3).

Braun and Clarke (2006) suggest two components within this phase, and each is now discussed in turn.

7.6.4.1. Extract-level review

This step requires the researcher to review all the *extracts* that have now been categorised under a theme; we must check and ensure that they do indeed form a clear pattern that we can group

7. Research Phase 4: Interviews

as a single theme. This step is effecting Patton's *internal homogeneity*.

Something I found difficult at this step was the breadth of some of the themes, and I worried that I had made the themes too broad. For example 'legislation', the 'power of incumbents' and 'sector capability' *are* all issues which will inhibit decentralisation/distribution (please refer to figure 7.2). Seen in overview, this has internal consistency. However, individually, all of these codes can be wildly different and could feasibly have branched off into a multitude of additional themes themselves. I was consciously avoiding doing this, since I was aware that the literature around RTA often condensed a *corpus* into fewer than five main themes²².

My candidate themes did appear, at this stage, to form a coherent pattern so I proceeded to the next step.

7.6.4.2. Corpus-level review

This step asks the researcher to zoom out and evaluate whether the initial thematic map "accurately reflects the meanings evident in the data set as a whole" (Braun & Clarke, 2006, p.91). Once again this is a normative determination and so, whether or not it seems 'accurate' is definitely subjective. The real question is – do the initial themes '*work*'?

At this stage, it became clear to me that having two themes which were essentially 'advantages' and 'disadvantages' of decentralisation was not analytically useful. In terms of RTA, at least, it seems rather facile to conclude that a major theme was "there are some advantages", or "there are some issues". Furthermore, it made no sense at the *corpus* level that I identified a theme – for example, 'access to data' – which could very well exist as a subtheme under either of these.

I therefore elected to completely review these two themes. Rather than my thematic map simply saying "some people thought these things were advantages", I really felt that I needed to pull out the key components and revise them as themes in their own right. What *are* the issues with decentralisation that people were talking about? In doing this I also had to revisit my initial fear of ending up with *too many* themes and this step therefore required a series of revisions.

In revision one I therefore decomposed the two problematic themes '*decentralisation issues*' and '*decentralisation opportunities*'. I did this by pulling the subthemes out into a list and attempting to either assign them to existing themes, or new and more meaningful themes. This is shown in table 7.6.

²²Indeed, course material made available on the Auckland University website by Professor Virginia Braun indicates that six themes would be a reasonable maximum for most projects (Braun, 2019).

²³Māori Data Sovereignty is not an area this writer can discuss, with any integrity.

²⁴This subtheme is a cynical benefit to the incumbent centralised hegemony.

²⁵This is a valid area for further research, but not a strong theme or subtheme in its own right.

7. Research Phase 4: Interviews

Table 7.6.: Thematic map: revision one

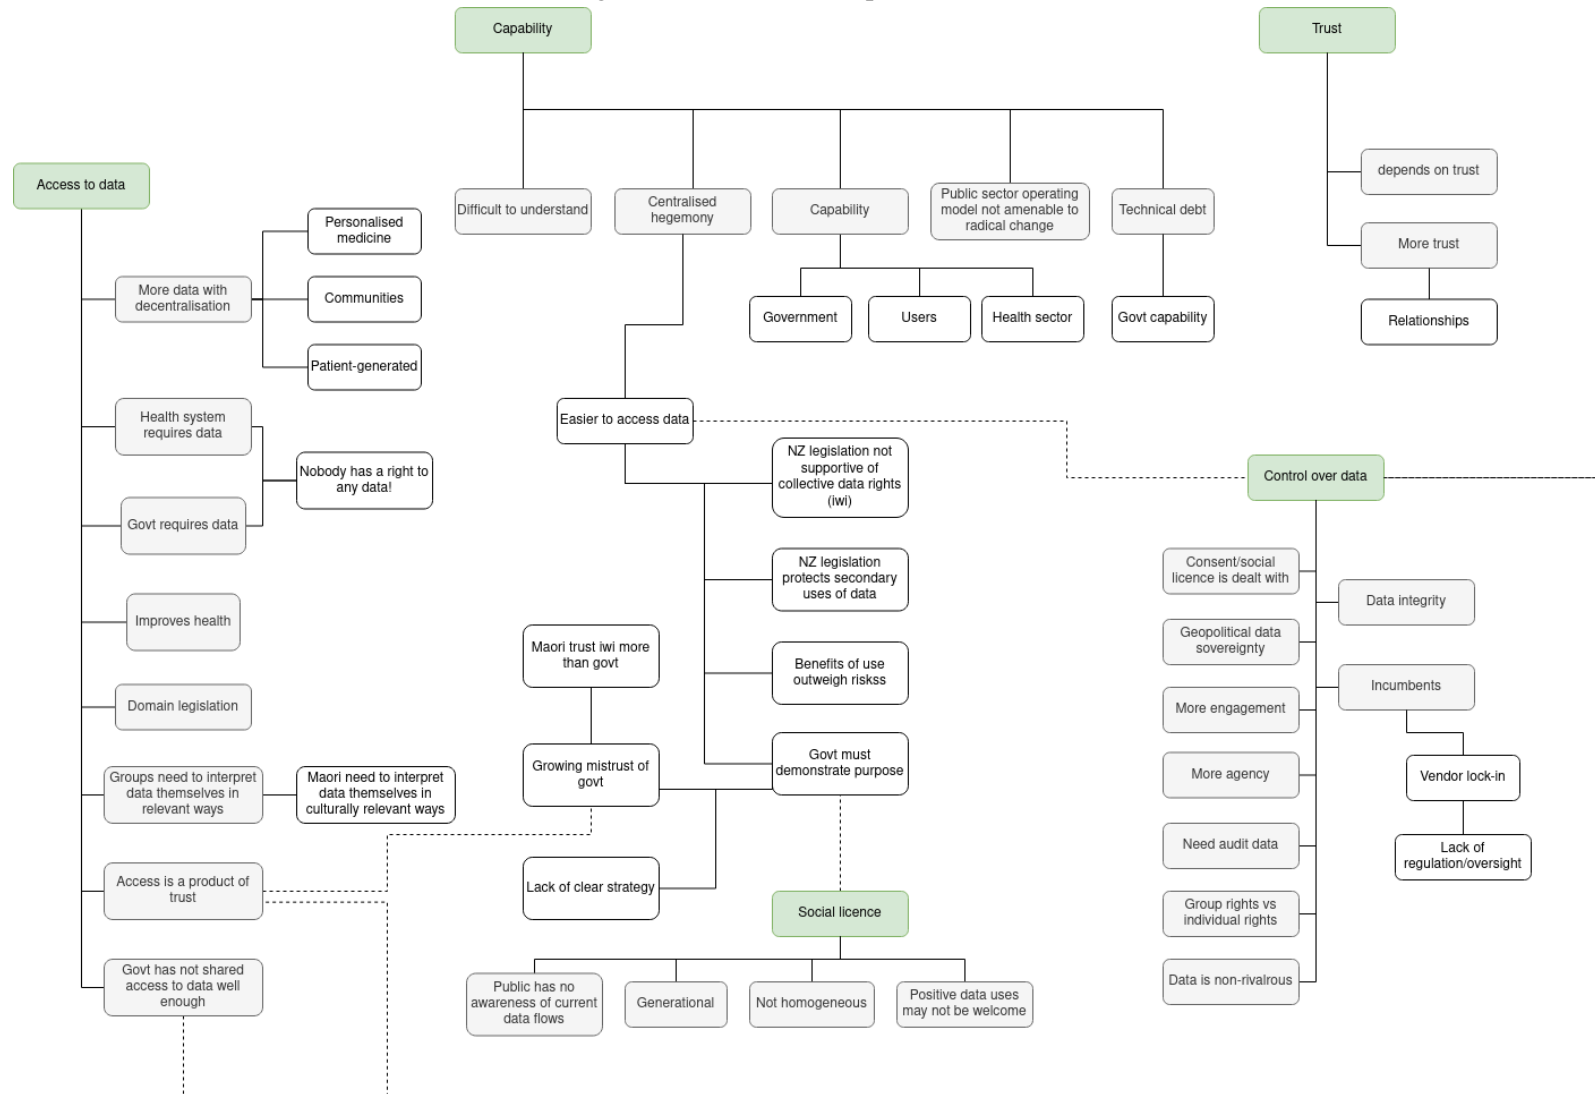
subtheme	Action
More access to data	Move to existing theme <i>Access to data</i>
Alignment with Māori Data Sovereignty	Remove ²³
Reduced risk/liability	Remove ²⁴
Data integrity	Move to existing theme <i>Control over data</i>
More trust	Move to new theme <i>Trust</i>
Difficult to understand	Move to new theme <i>Capability</i>
Centralised hegemony	Move to new theme <i>Capability</i>
Technical debt	Move to new theme <i>Capability</i>
Identity management	Remove ²⁵
Domain legislation	Move to existing theme <i>Access to data</i>
Dependent on trust	Move to new theme <i>Trust</i>
Capability	Move to new theme <i>Capability</i>
Incumbents	Move to existing theme <i>Control over data</i>
Public sector operating model	Move to new theme <i>Capability</i>

7. Research Phase 4: Interviews

This does now seem to highlight themes which are more meaningful. I recall from reviewing the transcripts, and undertaking the initial coding, that there were lots of comments around capability and trust. The revised themes therefore feel as if they are more reflective of the dataset with these now having a higher profile. The resultant thematic map can be seen visually in figure 7.3²⁶. It is a little more compact than the initial version, but the connections between themes and subthemes still make sense.

²⁶A full-resolution version of the image can be accessed at <https://imgur.com/a/HYtsqiZ>.

Figure 7.3.: Thematic map: revision one



7. Research Phase 4: Interviews

There is yet more refining to be done, however. Braun and Clarke advise that “you need to return to further reviewing and refining of your coding until you have devised a satisfactory thematic map” (2006, p.92). They also caution that this could go on forever and the researcher needs to recognise the point of diminishing returns in this iterative process and stop when it is *satisfactory*.

Reviewing the Thematic map: revision one, I had two main criticisms. Firstly, the new theme of *trust* felt disconnected and awkward, and secondly the thematic map overall needed to be simplified and streamlined (mainly in terms of subthemes). This led to a further round of refinement and in revision two I therefore took the following actions:

1. The *Social licence* theme was made a subtheme under *Trust*. While I had not drawn an explicit link between these themes in revision one, they do clearly relate to each other. If we recall that social licence is about the “permission ... [agencies have] ... to make decisions about the management and use of the public’s data” (Statistics New Zealand, 2018, p.3), then it will be obvious that this is a product of *trust*. The public will not permit management and use of their data if there is no trust²⁷.
2. The subtheme *Relationships are vital* was added back in under the theme of *Trust*. My renewed focus on the areas of *Trust* and *Social licence* led me to reconsider the importance of this theme. Specifically, reading the Statistics New Zealand survey on Social licence and trust, something that was prominent in the results was the fact that familiarity is positively correlated with trust. If people know about you and your work they are more likely to trust you; if they are not familiar with you then they will tend to have less trust (Statistics New Zealand, 2018). This seems rather obvious, in a way, but it did make me think again about the *Relationships* theme and how important that is in gaining trust.
3. The subthemes *More engagement* and *More agency* were merged into a new subtheme *Is empowering*. The original subthemes reflected the notion that having more control makes people engaged, and then motivated to make change or use the data positively. I felt this could be summarised by saying that having such control would be empowering.
4. The *Geopolitical data sovereignty* subtheme was removed. This code initially arose from people discussing Māori data sovereignty, and noting the confusion and ambiguity about whether its focus is on where the data is stored or the control people or communities can exert over that data. I have already noted in table 7.6 that this is not area I am qualified to

²⁷ Although do note several excerpts were coded as “Growing mistrust of government”, and yet there are few cases of individuals attempting to revoke permission. This is primarily because the dominance of centralised thinking offers no avenue to exert this, beyond the generic provisions of the Privacy Act 2020.

7. Research Phase 4: Interviews

address and, in many respects, this confusion is something of a moot point for this thesis given the prominence I have already given to concepts of *control* and *ownership*.

5. The *Need audit data* subtheme was removed. Originally this conveyed the need to see who was accessing and using data, and this has remained a very important component of the prototype. While it was prominent in the survey data we reviewed in chapter 4, it was not a strong theme from the interviews and does not add anything to the debate around how to implement decentralisation/distribution.
6. The *Data is non-rivalrous* subtheme was removed. This refers to the fact that data can be copied multiple times, and used by millions of people, without depleting its *supply*. This concept was raised during interviews in regard to trust and security. For example the concern that, when I share data with another user, it is impossible for me to control what they then do with that data. My rationale for removing this subtheme was predicated on reintroducing the idea of *relationships* (discussed in #2), since it seemed to me that the only reasonable approach to this issue is to consider human relationships as a basis for trust, even if a digital system is allowing you to make and manage those relationships. This also reflects some other excerpts, which asserted that the technical solution is only one piece of the puzzle. There is also a legislative component here, which is that the Aotearoa New Zealand Privacy Act 2020 is *purpose-based*. That is, rather than ascribing any concrete right to privacy (in the style of a Bill of Rights or Constitution), its primary concern is that information is used only in accordance with its original purpose. A decentralised/distributed system could quite easily allow the definition of purpose to sit closely alongside any data sharing agreement²⁸, to supplement the human relationship aspect. If unauthorised secondary uses of that data then occurred, falling outside the original purpose for which it was shared, it would fall under the provisions of the Privacy Act 2020 and the victim could seek remedy. This doesn't *stop* the unauthorised secondary use occurring²⁹, of course, but this is in fact no different to the thousands of data sharing relationships that occur today which 'suffer' from the same issue. Finally, I would add that any proposal to formally manage secondary uses of data in a digital system has two major issues of its own. Firstly, it reduces everything to a dependence on technical solutions which is not reflective of real life (and feels more aligned with a centralised paradigm). Secondly, the

²⁸In fact some enterprise data sharing solutions – such as Aotearoa New Zealand's Eight Wire (<https://eight-wire.com/>) – already implement this feature, although they are obviously working within a centralised paradigm.

²⁹In the same way that we cannot stop almost any other kind of offence or infringement, which are all prosecuted or remedied after the fact. Technology is shifting this, of course. For example, a self-driving car will only drive at the designated speed limit. But this capability is built on centralisation and 'enclosable carriers'.

7. Research Phase 4: Interviews

only viable technical solutions to achieve this are either moving back to formal centralisation, or utilising a blockchain³⁰.

7. The subtheme *Public has no awareness of current data flows* was removed, since this is already encompassed under the *Social licence* subtheme.
8. I merged the subthemes *Health system requires data* and *Govt requires data* into a new subtheme *Risks of not sharing data*. The two original subthemes were about the problems associated with central government, or health providers, not being able to access data because it has not been shared. The new theme draws this aspect out more clearly, whilst also allowing representation of the other side of this argument – that government may well require data to perform its functions, but it has done a terrible job of sharing it with others and using it effectively (Waitangi Tribunal, 2019).
9. The subtheme *Domain legislation* was originally intended to convey the concept that, whilst overarching legislation is supportive of decentralisation/distribution and data sovereignty³¹, specific domain legislation (for example, the New Zealand Health Act 1956 or the Medicines Act 1981) or regulation (for example, accreditation requirements of the Royal New Zealand College of General Practitioners) may not be. This subtheme was renamed *Legislation/regulation*, and other items in the thematic map were moved underneath it.
10. The subtheme *Groups need to interpret data themselves in relevant ways*, and related subsubtheme around Māori, was renamed as *Analytical/interpretive sovereignty*. This reflects the idea that sovereignty of data is not only about ownership and control. It also means that marginalised groups can conduct their own analyses and identify issues which are minimised or overlooked when such work is done at scale from the centre.
11. The subtheme *Access is a product of trust* was removed, and replaced with a connection being made between those two main themes to demonstrate the relationship.
12. The subtheme *Capability* is a replica of the main theme it belongs to. However this is expanded to note that there are particular capability issues across three groups: government, the health sector and end users. On reflection, I considered that this was encompassed in the existing subtheme *Decentralisation is difficult to understand*, and so this subtheme was removed. I consider this difficulty in understanding to be a product of another subtheme, *Centralised hegemony*, where the hegemony may constrain or compromise greater accep-

³⁰Which, I have argued in section §2.3, is centralisation in all but name and has other problems associated with it.

³¹Because it is centred on *purpose*, as we have already discussed in this section.

7. *Research Phase 4: Interviews*

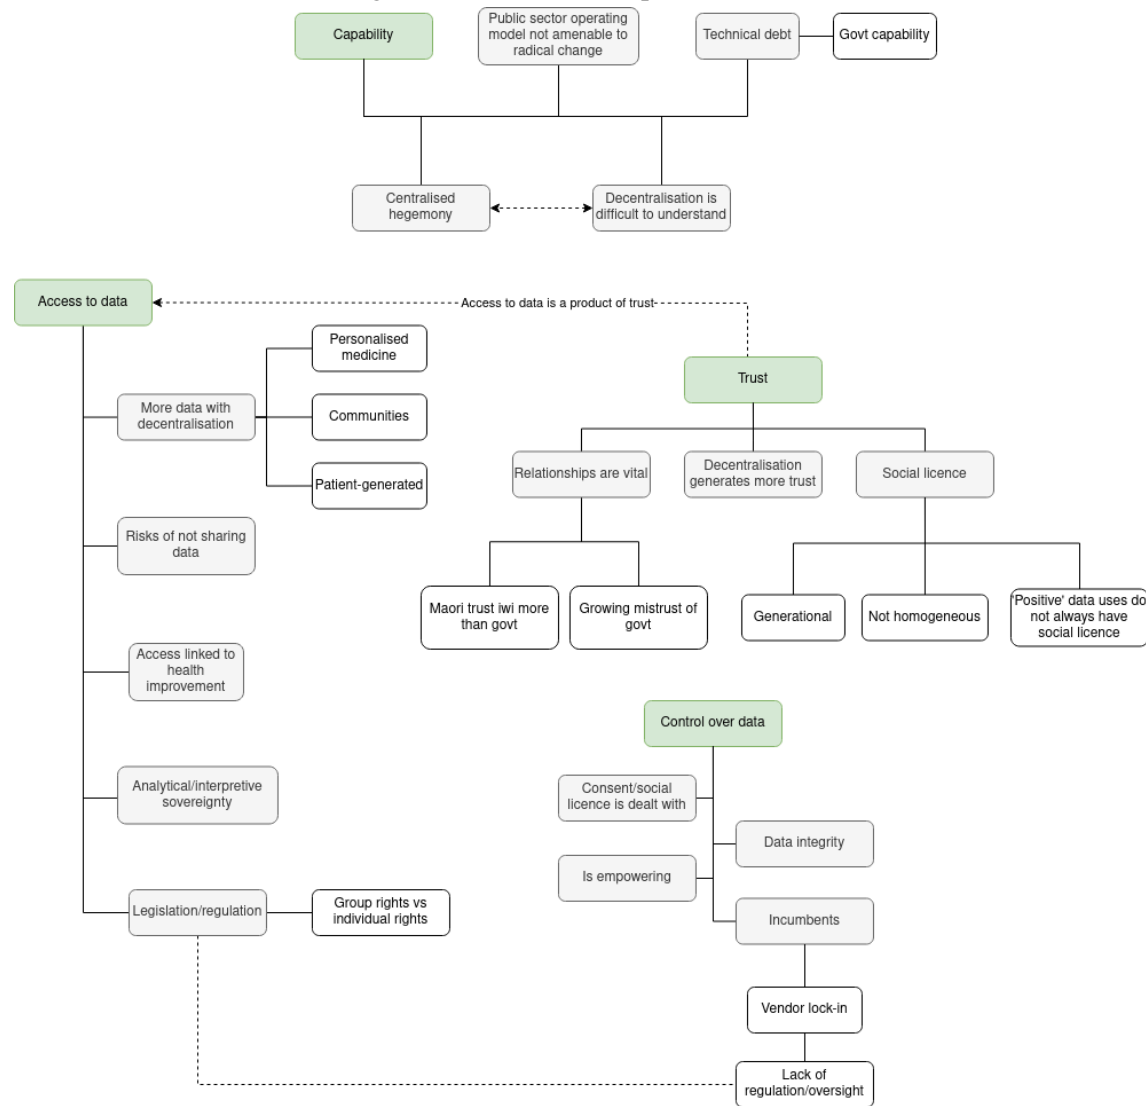
tance of decentralisation/distribution³². This concept is noted by joining the subthemes with a dashed line.

The revised thematic map from this last round of changes can be seen in figure 7.4³³.

³²One interviewee raised the issue of ‘status quo bias’ which is relevant here.

³³A full resolution version of this image can be accessed at <https://imgur.com/a/KuEPfXX>.

Figure 7.4.: Thematic map: revision two



7.6.5. Define and name themes

This is more of a refinement phase, where Braun and Clarke (2006) encourage researchers to ensure that themes properly identify the core message within each such that they are not too broad or complex. What is particularly interesting about this step (and RTA overall) is that we should definitely not be paraphrasing content; it is not a *descriptive* process. Instead we should be focusing on “what is of interest ... and why” Braun and Clarke (2006, p.92).

Having already modified some theme names in the previous phase, I was reasonably confident that the current version of the thematic map was coherent and consistent. I had managed to draw out four main themes, and used a range of subthemes to identify the points of interest. A test for this, as proposed by Braun and Clarke (2006), is to check that the scope and content of each theme can be described in two sentences. The results of this check are shown in table 7.7.

It was very useful to complete this check, and reassuring that I was able to describe these four large themes in a couple of sentences. As a result of thinking more descriptively about *Capability*, I also modified the thematic map to link the *Centralised hegemony* subtheme with the *Incumbent* subtheme in *Control over data*. This meant that all main themes were now connected to each other in some way, and provided further reassurance that there is an overarching story to be told from this data. The final thematic map is shown in figure 7.5³⁴.

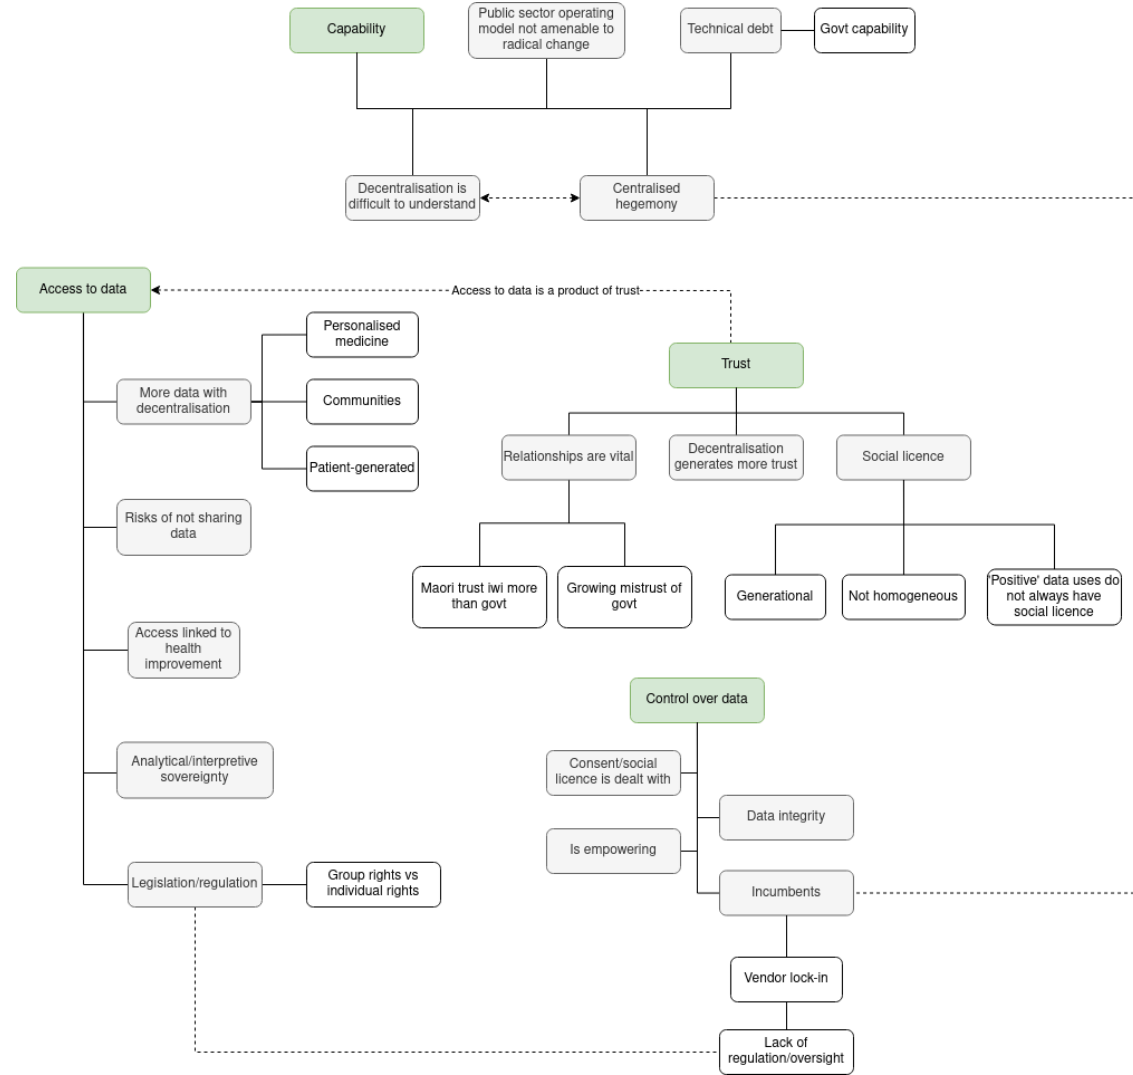
³⁴A full resolution version of this image can be accessed at <https://imgur.com/a/0cgMFqL>.

7. Research Phase 4: Interviews

Table 7.7.: Testing themes are clearly defined

Theme	Description
Access to data	The ability to have access to data, and the relative advantages it confers. This is mediated by a range of factors (legislation, business models, and technology).
Capability	The level of capability within the sector to shift to decentralisation. This focuses on the powerful technical and cultural hegemony of centralisation, but also deals with the perspective of individual users and clinicians.
Control over data	Deals with the perceived advantages of individuals being able to control their own data. Also notes the obstacle of incumbents (particularly in the vendor space), and their influence in maintaining the centralised hegemony.
Trust	The importance of trust – how this cannot be systematised, but yet it is the starting point for any sharing of data. Trust is highly variable and dynamic and I note as an important subtheme that there is growing mistrust in government, which has implications for the centralised foundations of government data collection and management.

Figure 7.5.: Final thematic map



7.7. Discussion

Having completed all preparatory steps, I can now start to analyse the data in earnest. The aim of this step is to provide a coherent, logical and interesting story from the data. Given the reflexivity embedded within RTA, I am not trying to uncover some latent positivist truth here – “the analysis of the material ... is a deliberate ... creation by the researcher, and must be constructed to persuade the reader of the plausibility of an argument” (Foster & Parker, 1995, p.204). Again, this should not merely be descriptive but should build an argument that informs the research questions. I will discuss each theme in turn.

Please note that each interview participant has been allotted a unique identifier, which will be used for attribution of any quotes in order to protect their privacy.

7.7.1. Access to data

In table 7.7 I described this theme as “the ability to have access to data, and the relative advantages it confers. This is mediated by a range of factors (legislation, business models, and technology).” There are four subthemes that arose.

7.7.1.1. Government monopoly on administrative data

The centralised paradigm is dominant and, under it, data is very difficult or impossible for individuals and groups to access. Pieces of data can be accessed, certainly, but there is no right to ongoing and comprehensive access to data. Participants outlined the problem in the following ways:

“How do we get data out of the IDI and into the hands of communities? This is sovereignty.” (Participant 1295cd16).

“There is no easy way to access that [centralised data] at the moment ... government still has a responsibility to share and collect and open data.” (Participant edc46a04).

“Data is locked in the DHB ... accessing basic aggregate public health data is really difficult... We shouldn’t even be in this position.” (Participant edc46a04).

In a centralised model, government agencies have access to all the administrative data and third party access to it is restricted. This is a product of the legislative context, where the core of the Privacy Act 2020 requires that data collected for one purpose must be used only for that purpose (with some exemptions). If the only bodies with the technical means and the legislative authority to collect data are government agencies (or parties to a statutory process such as healthcare

7. Research Phase 4: Interviews

providers), then they effectively *own* the consent process. This consent process is where purpose is defined – you permit use of your data in all of the ways described in that consent form.

The Privacy Act 2020 does include a right to access your data upon request (Principle Six), but there is no requirement to make it machine readable (as in the GDPR) and it is therefore essentially useless as a personal data source above and beyond intermittent review.

In the health context, this means that individuals have no access to their full health record³⁵ and yet the Ministry of Health can conduct any analysis they like on that data – as well as sharing it with other agencies, who may in turn share it with approved parties³⁶. It is certainly *possible* for community groups (iwi providers and other NGOs) to access this data, but there is an issue with capability. This is because such groups will be very infrequent users of that data and are unlikely to be aware of all the caveats and pitfalls which would be necessary for proper understanding. Of course, this is in turn a product of that data being constructed by government agencies and implementing ontological models which will not reflect the diversity, life experience and differing world views within the community at large (Morphy, 2016; Schultz & Rainie, 2014; Strongman, 2018).

Discussion around how non-government groups can gain access to administrative data begs the question “*who needs access to it?*” The answer is normative, in that we are accustomed to answer it from a current state perspective. The response to the question would perhaps reflect a status quo bias – “this is the way we’ve always done things”. The individual receives healthcare, the health system has the clinical data it needs and central government has the administrative data it needs. But individuals and communities are almost entirely sidelined from this data ‘partnership’ and this is reflected in the following excerpts:

“They [Māori] might want to do their own analysis, or augment it with some other data, and they can’t do that if they can’t get the data.” (Participant ca751e0d).

“From a Māori perspective, we want iwi and everyone to be collecting their own data, and we believe they should be accessing this data whenever they need to.” (Participant edc46a04).

³⁵Patient portals are attempting to fill this gap, but they can only expose what they know from their own particular silo of data. The most comprehensive offering within Aotearoa New Zealand, as at March 2022, is HealthOne (<https://healthone.org.nz/what-is-healthone>) but this is merely an additional layer of centralisation with another monolithic consent process.

³⁶An example of this is the Integrated Data Infrastructure managed by Statistics New Zealand, which takes in administrative data from most government bodies and makes it available to researchers. A total of 67 such research projects were carried out in calendar year 2021 – mostly by academic institutions (source: Statistics New Zealand, <https://cdm20045.contentdm.oclc.org/digital/collection/p20045coll17> accessed 20 March 2022). I am not aware of any research into public awareness of the IDI specifically, but it is reasonable to assume it is very low. Stats NZ’s research into their own social licence found that, while most people had *heard* of Stats NZ, more than two thirds “know just a little or almost nothing” (Statistics New Zealand, 2018, p.9) about their actual work.

7. Research Phase 4: Interviews

“Groups ... [should] ... have the ability to interpret and use their own data so they can put their own cultural framework on that data.” (Participant 1593596e).

These ideas go way beyond data sharing, or the building of more monolithic data silos to ‘share’ data with the rightful owner. They take us back to a first principles issue, around *who has the right to interpret meaning from this data*. The United Nations Office of the High Commissioner for Human Rights have presented a ‘Human-Rights Based Approach to Data’ (HRBAD), a key principle of which is *the right to self-identification*. The categorisation and use of identity data, according to this model, “should be developed through a participatory approach, to ensure respondents with these characteristics are optimally able to engage with the data” (United Nations Office of the High Commissioner for Human Rights, 2018, p.12). A useful example of this point is provided by another participant:

“Access to health records was a really significant issue for trans people saying they want to be treated as male or female. They didn’t want their health provider to know that they identified with a specific gender” (Participant df9eed6).

The same principle applies to any individual or group in the community, and I am asserting in this thesis that *everyone* needs access to their data (or, at least, they should have the ability to do so). The participant quotes above identify one of the reasons why this is important – groups should be able to conduct analysis from a specific worldview which may be marginalised. This is identified as the subtheme *Analytical/interpretive sovereignty* in the Final thematic map.

Going back to the earlier point around consent, we can also identify issues here around legislation. Two participants noted that the Privacy Act 2020 is technically supportive of decentralisation/distribution, since it is based on *purpose*. If a marginalised group wanted to collect personal data, using a decentralised/distributed approach, the purpose of data collection or sharing could be made clear in a way that met the needs of that group and the Privacy Act 2020 would easily facilitate this.

In a centralised system, however, this simply aggravates the ‘access to data’ problem since defining the purpose is left to those already in charge of the centralised system. In some cases this may be relatively benign. For example, MOH have certainly expressed a desire to share more data but do not yet have the technical foundation to do so. We might contrast this with incumbent vendors who are concerned with business models and revenue streams, and this is where a perceived lack of regulatory oversight has led to them having an undue influence on the flow of data across the health sector:

“If you hold the data, you’ve sort of been allowed to come to the point of view that you can control the data, even though I don’t think that’s an explicit setting.” (Participant 6f5c5542).

7. Research Phase 4: Interviews

Another meaningful issue with the Privacy Act 2020 is its failure to fully consider *group* rights around information and privacy. As one participant described:

“I think our law, and policy settings, framework that exists now, doesn’t recognise the collective very well. And so, iwi are a type of collective, and we have individual rights when it relates to data, and then we have rights and obligations of people who are working in the system. But a third-party collective that has the interests of their members is not well catered for by the current system.” (Participant 6f5c5542).

Indeed, the legislative purpose of the Privacy Act 2020 is to “promote and protect *individual* [emphasis added] privacy” (§3, Privacy Act 2020). It is true that group rights could be read in to the legislation, perhaps as a simple aggregation of individual rights. However, Kukutai and Cormack (2021) reject this and bemoan the global absence of collective privacy legislation³⁷. Nevertheless it seems odd, particularly in Aotearoa New Zealand’s ostensibly bicultural context, that the Act does not specify any provisions at all for collective rights.

7.7.1.2. The ongoing need for governments to access data

This was a prominent concept throughout the interview process – government’s *need* data to perform their basic functions. It is certainly hard to imagine government performing effectively without access to the administrative data it currently enjoys. But in the context of this thesis, this is not really a question of whether or not government should have access to people’s data so much as a question of how that is achieved.

If we imagine a fully decentralised health information system – *reductio ad absurdum* – then the data is owned by patients, who may decide not to share it with any government agencies. This would be an objectively bad outcome for the centre’s ability to plan, monitor and respond to health issues. But if we turn that same absurd scenario around, then there is the possibility for government agencies to have access to all the information they need, but with the explicit consent of every single individual which is a significant improvement on the current state. One participant remarked on this point in the following way:

“There potentially are some advantages ... in that if someone does give you access to that data, then you are very clear what kind of scope of use and what is appropriate to be done with that data.” (Participant 6f5c5542).

Neither of these two extremes is likely, at least in the medium term. It is hard to predict, of course, what people would choose to do given this choice and this furthermore depends on a

³⁷With the notable exception of Canada, where First Nations have developed and implemented a collective privacy framework called OCAP®.

7. Research Phase 4: Interviews

range of factors which I have not yet raised in this thesis – for example, the actual design and user interface of any system that people would be engaging with³⁸. So we have a dichotomy between the need to allow individuals to exert decision-making control over their data, and the need for government to continue accessing administrative data for planning and evaluation. Participants discussed this tension in the following ways:

“There is a role for the government and public health to play in helping people manage their health. And, in order for that system to work, the people who are trying to do population and public health need to have visibility of some of the data that might exist.” (Participant 4a626164).

“If you choose not to provide that data, then that significantly limits the Ministry of Health’s ability to respond. To actually do anything, because they don’t have any data there. And so, that’s a consequence of the individual exercising their own control.” (Participant 4a626164).

“One of the advantages ... from the way the system collects data now, is that for certain things, ... [it gets] data on everybody, or maybe data on everybody in a cohort. And so, that can be valuable when you are trying to understand, across the population, how is an intervention working.” (Participant 6f5c5542).

Many people would empathise with a government’s need for administrative data, yet we are left in a difficult position if we seriously want to consider decentralisation/distribution. The process from current state to ideal future state will likely have several stages, as I have demonstrated in table 7.8.

The precise order might change, of course, and government may wish to prioritise the social licence-building step well in advance. But this is merely to highlight that some concerted effort would have to be put in to building trust, so that government can continue to access the data it needs. This building of trust is something that could never be taken for granted again, since individuals would theoretically be able to withdraw consent to share data at any time. The ‘pay off’, so to speak, is a population who have understanding around the importance of government uses of data, who are willing to share it for specified purposes and – by virtue of having control over rich data of their choosing – can share broader and more meaningful data than government previously had access to³⁹.

There is another component to this, which can be described as the ‘operational efficiency of a centralised health system’. Under the centralised *status quo* (in Aotearoa New Zealand at least)

³⁸The user interface of any digital product is absolutely critical to how it is understood and used (Maier & Harr, 2020). Any modern app effectively lives or dies (in terms of popularity and uptake) on the user experience, and this is a broad, very specialist, field that is outside the scope of this thesis.

³⁹This last point, and stage 5 ‘Late decentralisation’, is discussed in further detail in section 7.7.1.4.

7. Research Phase 4: Interviews

Table 7.8.: Stages of decentralisation

Stage	Description
Current state	Government continues to access data it needs, based on a monolithic and rigid consent model from which individuals and marginalised communities are set apart.
Early decentralisation	Individuals begin to exercise sovereignty over data and government access is heavily reduced.
Social licence-building	A period of focussed effort from government agencies, with the goal of increasing trust and the rate of data sharing.
Steady state decentralisation	Individuals increasingly choose to share data with government.
Late decentralisation	Individuals and communities share <i>more</i> data with government than they had access to before.

it would be very wrong to think that we have achieved a health information utopia which must be protected at all costs from the threat of decentralisation/distribution. Even with highly centralised systems across the health sector, we are still as far away from realising interoperability as ever. In fact this centralisation may very well be an *active barrier* to interoperability. One participant commented thus:

“My own experience ... involved being a patient across two DHBs. And it was terrible. In fact I did end up having to take control of my data because otherwise hospitals simply wouldn’t share my data with the other DHB.” (Participant ca31cb2c).

In a centrally funded public health system as relatively small as Aotearoa New Zealand, it is indeed a curiosity why information sharing between two neighbouring DHBs is still such a problem today. This issue has been the centrepiece of health information strategies for decades. Gauld surveyed the troubled history of health information management in Aotearoa New Zealand in 2004, concluding that good intentions were stymied by the ‘New Public Management’ of the 1990s; the purchaser/provider split actually *encouraged* data silos and provided no incentive to think about sharing and distribution in design (Gauld, 2004). My own research into this area found that, by 2016, nothing meaningful had changed even though a new statutory body (the New Zealand Health IT Board) had been setup in 2018 with interoperability as a goal. Additional findings from this research were that there had been a trend of groups taking up strong incumbent positions which would be difficult to shift and that there was verifiable mistrust even amongst different stakeholders within the Aotearoa New Zealand health system (Poor, 2016).

7. Research Phase 4: Interviews

A decentralisation/distribution apologist would interpret all this to simply be the inevitable outcome of centralisation, both in technical architecture and mindset. Nevertheless, the point I wish to make here is simply that the system we might be afraid of impacting in a move towards decentralisation/distribution is not necessarily one to lionise; and that the groups who express anxiety about it will only be those with something to ‘lose’.

7.7.1.3. Risks associated with lack of access to healthcare data

Finally, perhaps a subset of the previous subtheme, there are specific professional groups who require access to specific data within the state sector. In our case, government access to data for planning and evaluation is one thing. But the potential for health data to be withheld from health professionals working with an individual or whānau is unthinkable.

In the case of the government, the perceived need for personal data is perhaps more distant and abstract. We know they do things with data, and we assume it is used responsibly and effectively. Our certainty around this is a product of the building of social licence, which has been sporadic thus far. In the case of health professionals, however, it is very different:

“If you want your doctor to be able to make the best decision that they can about your health situation, then they probably need to have as complete a picture as possible.” (Participant 4a626164).

“If you get into a situation where an individual can basically hide information that they don’t want their doctor to see, then that can actually compromise health outcomes as well.” (Participant 4a626164).

We often have a direct and ongoing relationship with health professionals, and they are some of the most trusted people in society. In fact, Ipsos reported in 2021 that Doctors are now the most trusted profession globally (Clemence, 2021)⁴⁰. Indeed, in section §4.9, we saw that survey respondents would be very happy to share their personal data with their GP, but less happy generally to share it with government agencies.

From the evidence presented in chapter 4 we should be confident that, in a decentralised/distributed system, patients would be sharing their data appropriately with health professionals who are caring for them. It may be that a solution design for such a system implements some kind of ‘break glass’ facility, such that health professionals would always be able to access critical information (such as patient allergies, or adverse reaction history, when prescribing medication). But it is very clear that health professionals would be unable to perform their duties without access to relevant personal data, and this should be a central component of any RDHIS.

⁴⁰Politicians and government ministers were amongst the least trustworthy professions, which indicates an issue for the building of social licence in a decentralised/distributed system.

7.7.1.4. Decentralisation opens the door to accessing *more* data

The subthemes so far in this section have all expressed some anxiety about potential negative consequences of decentralisation. However, there is a pathway here for government and other stakeholders to have access to more data than they ever did before. Such access would need to be permissioned by the individuals concerned, to be sure, but no such pathway exists for government – or indeed anyone – under a centralised paradigm:

“There is hugely valuable information that sits in communities that will never ever sit in the administrative data that government collects and holds. So what are the conditions, what are the relationships that you need where communities are prepared to share that information?” (Participant 342775a4).

“It’s a partnership – so government are sharing, and communities are sharing and we get the best of both worlds.” (Participant 342775a4).

What the above participant quotes indicate is that there is a rich, untapped, source of data sitting in communities that would provide valuable intelligence to government that is completely absent from the administrative data at present. This may sound idealistic, but what I am calling ‘Late decentralisation’ is a stage where the ceding of sovereignty to people over their health data, and a concerted programme of social licence-building activity, leads to a high-trust environment where people see the impact of their data in meaningful ways – they see *themselves* in data and policy-making, and want to contribute more.

This will certainly not be easy, but the potential is to have fully engaged citizens willingly share their data with government agencies, rather than the disenfranchising and marginalising scenario we see currently. One of the risks of this admittedly rosy outlook, is that the generation of social licence is uneven. For example, marginalised groups may take longer to come to that position of trust. They may never get there, in fact. And, in a scenario where they have the ability to withdraw their data, we may see the deleterious scenario where marginalised groups become *more* absent from analysis and policy-making than they were to begin with. This is indeed a wicked public policy problem, borne of pre-existing mistrust which I attribute to the flow-on effects of centralisation.

7.7.2. Capability

Capability here encompasses a few different dimensions. Because decentralisation/distribution is such a new concept, it poses challenges to all stakeholders – from IT specialists, to business users and individual citizens. Capability means different things to all these groups, but it is something that is essential to understand in terms of implementation.

7. Research Phase 4: Interviews

7.7.2.1. State sector

Participants commented on state sector capability in a number of ways, most of it negative:

“When somebody snaps their fingers and says “set that up” and they have to set it up in 24 hours, the last thing they think about is how they might do the data piece. And in fact, even when they had more than 24 hours, the last thing they thought about was how they would do the data.” (Participant 342775a4).

“In government, you’ve got endless resources to build whatever you want. No one there has to think about where their money is coming from. In the corporate world, if you are not demonstrating value to your customers, you don’t get paid. The approach in the corporate world is to go and try something small and test it and add value, rather than the government approach to just spend millions of dollars on big things without stepping back and thinking about the outcome they’re seeking.” (Participant 1295cd16).

“There is a high degree of variability tending towards the low end of maturity where we wouldn’t expect [governance bodies] to know what privacy provision should be in place, what access control should look like. What you end up with, in a sense, is the worst of all worlds.” (Participant 682d3989).

“Most people in government don’t have the skill or capability to engage in a good discussion about that [decentralisation]. You just get so focused on the thing that’s in front of you, and people don’t have enough time to step back and picture what they want Aotearoa to look like in the future. It’s not hard, it’s just people being bold enough to have some strategic conversations and long term aspirational conversations.” (Participant 1295cd16).

“When I have an interaction with the health system, I know my data’s going to be treated poorly – not with ill intent, but just because of clumsiness.” (Participant ca31cb2c).

There is a sense, amongst these comments, that government is simply not geared to deliver the radical kind of change implied in a move to decentralisation/distribution. Whilst this may be seen negatively, it is a perfectly sensible product of the democratic process where change has to be deliberated, reviewed and agreed, before other levers and instruments are put in place to implement it. This is harder again in the technology space, where change is even more rapid – “given the Fourth Industrial Revolution’s extraordinarily fast technological and social change, relying only on government legislation and incentives to ensure the right outcomes is ill-advised.

7. Research Phase 4: Interviews

These are likely to be out-of-date or redundant by the time they are implemented” (World Economic Forum, 2016, p.6).

Another theme in the above excerpts is something akin to a lack of *purpose* or *deliberation* in government activity; a notion that government is unable to consider the ‘big picture’ and systematically plan for defined outcomes. This is hard to quantify, of course, and seems like it might be an easy thing to say when standing outside of government. But it is a notion that felt prevalent to me during the interview process and feels familiar from my own work within the public sector. Moore interviewed Aotearoa New Zealand social service providers in their 2019 study of the impact of ‘social investment’, and noted similar and very relevant sentiments from participants, one of which I will reproduce here:

“A lot more data was collected. And there was a sense that they hadn’t figured out how to do it either. But they knew they had to deliver something. And you know, you would have an audit ... You would gather everything they asked for last time and you would have it all in front of you and they would say “well, actually we want a measurement of this.” And no one has asked you for that data before. And although we have a client management system, that is not something that we have asked it to gather so we are not entering that data. ... So there is this sense that a lot of thought and a lot of change has gone in but it hasn’t actually filtered into the workers on the ground in MSD and Oranga Tamariki (as it is now) in terms of what that actually looks like in a tidy way, that is my experience”. [Community Provider 4] (Moore, 2019, p.133).

So we cannot rely on government to articulate a legislative/regulatory landscape that would facilitate a radical shift towards a RDHIS (or perhaps even anything). This implies more of a ‘bottom up’, or even ‘middle out’, approach would be needed (Coiera, 2009) and to some extent that is what we’re seeing with the MOH *Hira* programme which will “enable access to a virtual electronic health record by drawing together a person’s latest health data from trusted sources” (Ministry of Health, 2021a, para. 2). To be clear, MOH are leading this work programme but extensively involving stakeholders in design and approach with the aim of creating a *platform* – rather than a product – upon which third parties can build. It aims to achieve many of the RDHIS goals, but firmly within a centralised paradigm, because that is where we are at present. This was summarised to me by one participant in the following way:

“I think there are things about the decentralised model and what it’s trying to achieve and the benefits of it that we want to aspire to, even if we don’t necessarily achieve it in a fully decentralised way” (Participant ca751e0d).

7. Research Phase 4: Interviews

This begs the question – if we can achieve the key facets of a RDHIS in a centralised way, then what is the point of a RDHIS? The answer to this question goes back to ideas around power and control, analytical sovereignty, political neutrality and the ascribing of full – rather than merely curated – data sovereignty to people and groups.

One other component raised by interview participants is that of ‘technical debt’. This is a commonly understood term in computing and solutions architecture, but can be comprehensively described as follows:

“In cities, repairs on infrastructure are often delayed and incremental changes are made rather than bold ones. So it is again in software-intensive systems. Users suffer the consequences of capricious complexity, delayed improvements, and insufficient incremental change; the developers who evolve such systems suffer the slings and arrows of never being able to write quality code because they are always trying to catch up” (Booch, 2015, p.ix).

The concept is generally about short cuts or compromises that are made across a wide range of dimensions but, crucially, are not repaid or dealt with properly. You do not need to have experience in delivering public sector IT projects to imagine that our health sector does indeed have considerable technical debt. It will not add much to this part of the discussion to examine some of the issues we see in Health IT, particularly when senior government officials are candid enough to say:

“In health anyway we’re talking about quite significant technical debt. The ability for those systems to participate in that model would be basically impossible, or there’d be significant investment required to come up with some means of doing it.”
(Participant ca751e0d).

This certainly does contextualise the scope of the *Hira* project, but makes one wonder whether or not – if the foundation is really so fragile – it is a suitable one on which to build any new kind of platform⁴¹. But the extent of this technical debt represents a practical issue for government, in being able to pivot to new approaches. Health systems are obviously *complex* systems (at least in terms of data movement), and there is an element of always having to design solutions for the lowest common denominator. When 96 GP Practices are still not using the Aotearoa New Zealand ePrescription Service (according to data available at the time of writing in May 2022 (Ministry of Health, 2022)) we must assume they are continuing to hand-write or fax prescriptions. Even if all pharmacies are ready to accept electronic prescriptions, they must

⁴¹But of course, *something* has to be done and what is the best approach which will cause minimum disruption to the functioning of the extant health system? These are extremely complex issues which should not be underestimated by onlookers or commentators.

7. Research Phase 4: Interviews

remain able to receive faxes and process hand-written prescriptions. This is one example of the long tail of implementation which contributes to technical debt and constrains system agility.

7.7.2.2. General public

Interview participants (almost every one, in fact) raised concern about the capability of the general public in negotiating a RDHIS. This seemed to focus on the perceived burden of managing the sharing of data, or communities/groups becoming competent with generating and utilising their own data. These could be grouped into a type of ‘digital literacy’ issue:

“We need to support people to collect their own information. Marginalised communities shouldn’t be expected to do all that on their own with very little knowledge on how to do that.” (Participant edc46a04).

“The idea that you can be interactive with your health information – it presupposes lots of layers of skills and knowledge and access to the internet and information and so on that I think make it quite problematic for some groups of people.” (Participant df9eed6).

“To understand the rights, rules, obligations [of a decentralised system], and to give them confidence with that, that’s really hard stuff. This is not an easy quick fix.” (Participant 342775a4).

“If we were moving to a decentralised world tomorrow, so much education will be needed but people don’t get it because they’ve been living in a centralised world for so long... It’s work that will not be achieved through one-to-many advertising. You’re not going to do TV ads and leaflet drops and suddenly everyone’s going to understand this new world of decentralised data. It requires a much deeper level of engagement in order to get people to understand what this now means for them.” (Participant 4a626164).

“We operate with a wide spectrum of levels of sophistication of people who are either providing the services or accessing the services. So, how difficult is it to retain and control your own records? Is that something that everybody could do, or is it something that only the more tech-savvy members of the community would be capable of?” (Participant 6f5c5542).

The underlying concept here is that centralisation has disenfranchised people from being able to think critically about data and its usage. We have already identified some of the issues around capability in section 2.2.3, where we reviewed literature suggesting that individuals have become fatalistic about uses of their data or their ability to control it (McMullan, 2015; Solon, 2018).

7. Research Phase 4: Interviews

The nature of the data landscape, predominantly (but not only) in the private sector, has meant that individuals have become de-skilled around critical issues such as consent and transparency (Data Futures Partnership, 2017a; Malkin et al., 2018). Whilst I would argue that centralisation has been a springboard towards this situation, technology does not exist in a vacuum and it is perfectly feasible to do consent and transparency well even in a centralised model. Why hasn't this happened? The hubris of researchers interviewed by Jao et al. is instructive – “broad consent can sufficiently support autonomy in informed consent by acting as a decision to allow *others to decide*” (2015, p.273)⁴².

I accept that that quote could be interpreted in different ways. My negative interpretation as *hubris* is based on a follow-up quote by a researcher who minimised the potential impact of a privacy incident on participants when using this model, as well as the underlying attempt to justify broad consent models whose chief benefit is to make the research task easier for the researcher. Having said that, this notion of delegating consent is extremely pertinent to this perceived capability issue and points to one way that the burden of managing permissions might be mitigated. I should say, before going further, that we are now firmly in the territory of designing how an end-user decentralised app might look or operate. This was meant to be out of scope for the prototype design, but cannot be ignored when discussing real world use cases. On this basis, anything is possible – an optimal design and user workflow would require the expertise of specialists in the fields of User Experience and Human-Centred Design.

One option may be to allow ‘data sharing relationships’, which provide the ability to assign high levels of trust to groups or individuals who then have the capability to share further on that user’s behalf. For example, I may decide to give ‘Full Admin’ permissions to my GP or specialist. She is then free to share my information further, as required. Because the system will keep a full audit trail, I will still be able to see all of that secondary and tertiary sharing and access taking place⁴³. Alternatively, I may really prefer to manage and configure bilateral sharing arrangements in great detail. The point is that people should be able to manage access to their data in a way that feels right to them. Striking this balance is firmly in the design and usability space; it can easily be accomplished in the technical layer.

A corollary of this concern about capability and digital literacy is a concern around equity. One participant described it thus:

“If the amount that you are prepared to engage with it [a decentralised system] affects the benefits that you get from it, then that could be a problem because we already know that equity is a problem in the current system.” (Participant 6f5c5542).

⁴²Although this is not by any means intended to single out researchers; simply that this attitude is likely to be one that has aggravated the impact of centralisation on user education around consent and transparency.

⁴³And, crucially, will be able to repeal any permission if I so wish.

7. Research Phase 4: Interviews

So, following this logic, confidence with using a RDHIS to your (or your group or community's) advantage is contingent on underlying digital capability and access to digital services. If it is socially marginalised groups who will be *least* capable of accessing and using this system, then none of the benefits will accrue to them and this is likely to compound existing inequity. Similarly, we have already noted in section 7.7.1 that 'analytical sovereignty' is one way by which marginalised groups can raise the profile of their lived experience. However, this is dependent on being empowered to do so from a knowledge perspective – as well as more practical issues, such as access to devices and internet provision. The Aotearoa New Zealand health system must cater for all New Zealanders and, even in 2022, this means that a blanket reliance on digital services cannot work. It doesn't mean we cannot think and plan about using technology to advance wellbeing, but we must pay equal attention to those we are potentially excluding from our grandiose proposals.

7.7.3. Control over data

Control over data must be considered very separately from simple *access* to data. I noted a trend towards conflating the two during the interview phase which I, perhaps rather unkindly, marked as a 'capability' issue. I must admit I am not entirely sure it is the case that people always tend to think of them together, or default to thinking about access simply because the concept of control seems so absurd when we have been socialised into a centralised mind set. But I do view control as being really central to the notion of 'sovereignty'.

Access to data can be a very passive thing. The controller of a data store might provide you full access to all of your data – even, perhaps, in a fancy app. But you don't have any direct *decision-making capability* over that data. It might be siphoned off to another party and, even if you were permitted to see that happening, you would not be able to stop it.

To be sure, we can actually get reasonably far in a centralised paradigm, but it takes a great deal of effort to purposefully build in the functionality and auditability that at least minimises the impact of not having sovereignty over that data. So, 'control over data' means that the individual (or group, or community) has immutable decision-making authority over their data.

Because we are so fixed in a centralised paradigm, this can be a difficult and abstract thing to think about. Indeed, it was a key theme from the interview phase and there were a number of quite varying subthemes which arose.

7.7.3.1. Consent and social licence

There has been much thinking and writing about consent and social licence, some of which we have already reviewed (Aitken et al., 2016; Data Futures Partnership, 2017a; Statistics New

7. Research Phase 4: Interviews

Zealand, 2018). Broadly speaking, the literature is concerned with understanding the extent of social licence held by different bodies, and how to increase that level of social licence. It should be obvious that this line of thought is very much borne of centralisation. Why would we even need to think about any of this, unless there is already some perceived imbalance of power or concern about how individuals' data are used? The potential promise of a RDHIS is that consent effectively disappears as a meaningful issue:

“If someone does give you access to that data [under decentralisation], then you are very clear what kind of scope of use and what is appropriate to be done with that data, because that has been presumably negotiated through the transaction. And so there wouldn't be the same uncertainty that perhaps there is now about, you know, are we able to do this legally, do we have social licence to do this, how do we weigh up whether or not we should be doing this with data that we just happen to have? For an agency it's a nice thing to have.” (Participant 6f5c5542).

In this scenario, consent has been defined and actioned by the owner of the data. There is potentially no need (although, again, this depends to a large extent on the design of the app/interface, and the digital capability of the user) to ask any further questions around purpose or even to further review or monitor consent. When the user wants to withdraw that consent, they can easily rescind any access to practitioners, groups of practitioners or organisations⁴⁴. This potentially has subsequent benefits in governance of data projects which utilise personal information – making the task of Privacy Impact Assessment or research ethics somewhat simpler.

Social licence is a different proposition, since it is the thing that will determine whether that consent transaction is made. In a centralised model we tend to see the two things as being directly related, or even conflated. In section 7.7.1 we noted how the Privacy Act's focus on *purpose* assigns disproportionate power to those who are already in charge of the centralised system by allowing them to define purpose. The centralised data store controller⁴⁵ designs a consent form which specifies the purpose for which the data can be used or shared. The end users' decision-making capability within this process is restricted to simply consenting or not consenting.

The purposes in a consent form may be very well described, or very general – leading to some possibilities for misinterpretation or *over-use*. For example, if my information is to be shared with “a health professional involved in my care”, I am likely to be comfortable that this encompasses my hospital specialist or my community extended care team. But this wording

⁴⁴The potential groupings of permissions, and the need for identity frameworks, is discussed further in chapter 8.

⁴⁵Who, in the government setting, may be an agency governance group, a senior government official or the Privacy Commissioner. In a primary care setting, it might be the practice manager or the GP owner.

7. Research Phase 4: Interviews

could also encompass my physiotherapist or community pharmacist, with whom I would prefer a more limited data sharing relationship.

So in this case, consent can be very passive for health service users and is often literally a ‘tick box exercise’. Acknowledging this scenario, and the ambiguity of some consent processes, agencies can spend considerable effort anguishing about the trust they may or may not have to use data in a particular way. The consent horse has already bolted, and yet we still have to consider *ex post* the social licence we have in order to utilise the data. In reality these should be very distinct activities.

In a decentralised/distributed paradigm these components are indeed separated. A user trusts an entity and their ability to safeguard their data; the user actively consents by providing that entity with access to their data on whatever conditions they wish (for example, it might be only subsets of data that are shared or the sharing might be time-limited). It is very clear to all parties what consent that user has given, and for what *purpose*. The user is able to define how the Privacy Act will be used to protect them, in the event of a breach⁴⁶. The entity only needs to worry about social licence where meaningful numbers of people are choosing not to share their data. In this way, the path to data sharing is controlled by the user and is sequential and very clear.

The distinction between these two paradigms is shown in figure 7.6 and figure 7.7. In the former, albeit highly simplified, we can see that Purpose is the largest object. This reflects the ability for the data store controller to define and expand this as they wish⁴⁷. The consenting user may not be fully cognisant of the potential reach of that purpose at the point that consent is given. In fact, the user may have reservations about what they do understand but, because consent is usually all or nothing, they cannot exert any real decision-making power over this process. Social licence (or the agencies own perception of it) will mediate the Purpose, to some extent at least. As before, the user may or may not attribute social licence to the data store controller, but consent is not flexible enough to allow any nuance. The user is relatively passive.

The matrix behind the image shows nodes which represent potential data uses – many of which are within the Purpose sphere, but not within the Consent sphere, demonstrating that users will be unaware of intended or actual purposes. Only one node is shown at the conjunction of all three spheres. This is, of course, a highly subjective representation, but I am merely trying to demonstrate the ambiguity of the *status quo*.

Figure 7.7, by contrast, shows a much clearer and circular relationship between the core con-

⁴⁶This is once again veering into app design, but it would be trivially easy to allow the user to specify allowed uses of data, as part of a data sharing agreement, by which the receiving entity would be obliged to abide. The audit facility would provide a mechanism to check compliance with the agreement.

⁴⁷In the case of health data, any such expansion is obviously within relevant legislation and regulation. But enough flexibility remains to result in ambiguity and consternation, as seen in the WOCA example (Neilson, 2021).

7. Research Phase 4: Interviews

Figure 7.6.: Consent and social licence conflation under centralisation

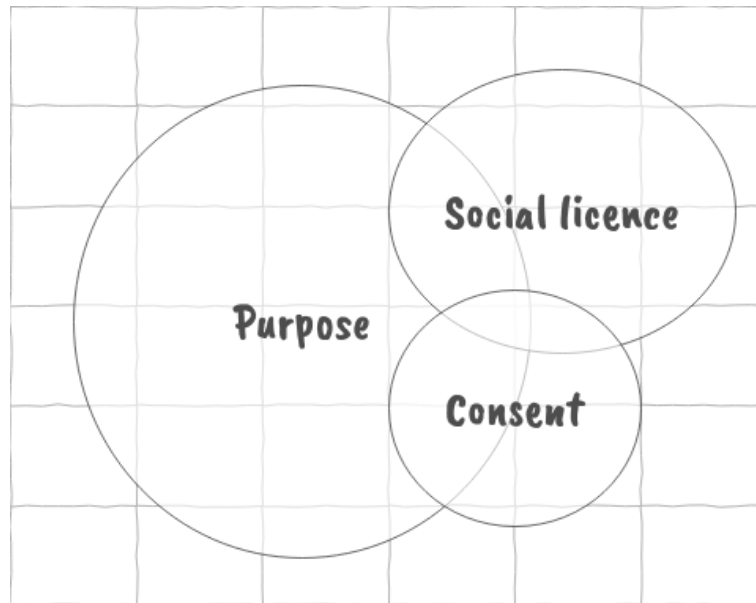
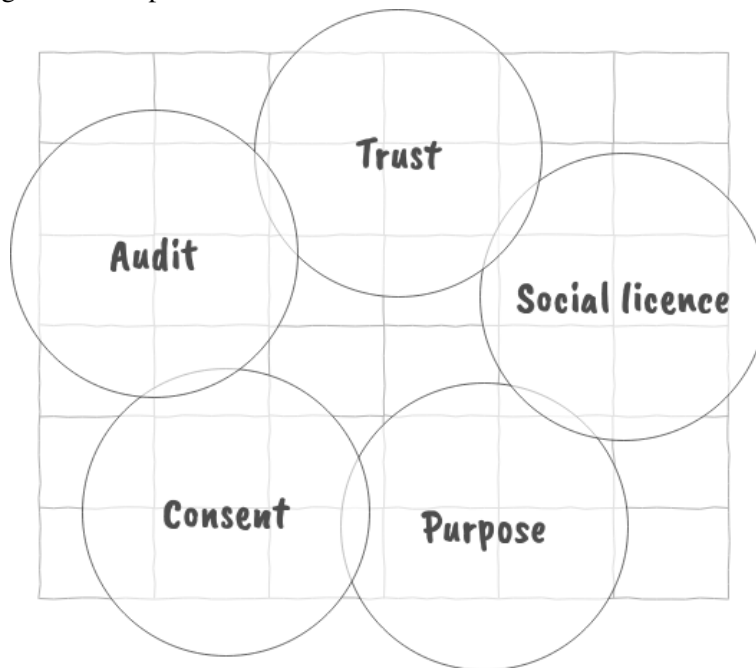


Figure 7.7.: Separation of consent and social licence under distribution



7. Research Phase 4: Interviews

cepts. In this representation, Social licence is a product of Trust. If I trust MOH generally, then I will feel more confident in their ability to use my data appropriately. The level of Social licence defines the Purpose with which I choose to share any data, and the sharing of that data provides a clear and explicit consent to the party I am sharing with. Finally, I am able to review how my shared data is being used and, if the terms of the data sharing are being met, this will further increase my Trust in that party. Once again, this is a highly subjective representation, but it does serve to show a potential ‘virtuous circle’ of trust and data sharing where Purpose is defined (or at least very clearly agreed to) by the user, and this can be monitored.

Given some of the ambiguity and conflation that we have discussed above, some participants noted that, overall, it is much better to have the clarity of decentralisation/distribution from a risk perspective:

“This is one of the advantages of the decentralised approach, is that it takes the risk away from the central agencies of having responsibility for governing and protecting that data. Because, the more [that is held], the more the risk is that someone will do something inappropriate, or ... use that data in a way that is inappropriate or not as intended.” (Participant ca751e0d).

So, whilst we acknowledge the danger of low social licence and lower than anticipated data sharing, we also see that perceived risks of having to manage large centralised data sets are *removed* in a RDHIS. These risks span multiple domains – security, privacy, governance – and can absorb a large amount of time and effort. It also highlights another social licence risk, which is where a purpose is ambiguous and a holding agency are asked by a *third party* for access to data.

This eventuality came to pass in late 2021, when the Whānau Ora Commissioning Agency (WOCA) requested detailed data from MOH around vaccination status for Māori. The MOH held all this data, and was being requested to provide it to WOCA so that they could follow up and work to increase vaccination coverage rates. MOH were put in a difficult position as custodians of that data because the Privacy Act 2020 was somewhat ambiguous in this case and, at the same time, many Iwi (including Ngāi Tahu), were objecting to release of that data to WOCA (Neilson, 2021).

At the risk of oversimplifying to support my own argument, this situation would not have occurred under a RDHIS; people could share their data with WOCA if they chose to. If WOCA were concerned that the people they needed to reach were not sharing their data, then it would be incumbent upon them to work with iwi, or community groups, to build social licence. There is absolutely no ambiguity, or possibilities for differing interpretation, in this scenario.

7.7.3.2. Data integrity

Data integrity is a concept often used interchangeably with Data quality, but they are in fact distinct concepts which are worth unpacking. To begin with, data quality is a component of data integrity. You can have data quality without data integrity, but not *vice versa*.

Data quality is concerned with “the usefulness, accuracy, and correctness of data for its application” (NIST Big Data Public Working Group Definitions and Taxonomies Subgroup, 2019, p.11). It is not an objective standalone concept, since ‘usefulness’ means different things to different groups and, indeed, will change over time. This touches on some of the concerns raised by participants (and discussed in section 7.7.1) around the overall quality of data they might have access to, in the event that people were choosing not to share with government agencies:

“If you need, for the particular thing that you’re doing, a really large sample, can you get enough people? If you need a representative sample, can you get enough people in all of the different groups that you would need to get?” (Participant 6f5c5542).

So there is certainly a clear data quality issue here, in the event that MOH need to evaluate a programme or policy implementation with only 50% of peoples data – which may, in turn, be disproportionately skewed towards one demographic or another.

Data integrity is a higher order issue, more concerned with ensuring that there has been no *unintended* change to the data at any stage of its lifecycle. A common example of data integrity in action would be how incoming updates to data are handled by a receiving database. The receiving database must ensure several things occur:

- The data is linked with the correct record
- That the updated components are updated
- That no other components are updated.

Data integrity faults in such a scenario might see the creation of duplicate data, or incoming data being stored against the wrong record. In any example, some data has arrived but it has been misinterpreted or processed incorrectly, such that it is not accurate in the receiving data store.

Participants commented on quality and integrity issues in the following ways:

“One of the things [issues with decentralisation] that’s coming to mind is how do you maintain the integrity of the data” (Participant 1593596e).

“We will get quality data because it is managed by that person” (Participant 682d3989).

7. Research Phase 4: Interviews

“The outcomes that we are seeking is higher quality data that captures more than an administrative process” (Participant 1295cd16).

“The other benefit of decentralised data, which is that there’s only one copy of it. I’m not trying to maintain the integrity of my copy of it ... People at the moment ... are fixing data errors at the point of ... capturing and storing it, not going back to the source and fixing it at that that level. So you’re just instantly creating an integrity issue” (Participant ca751e0d).

The first two quotes can almost be viewed as a direct question-response pair (although one talks about integrity and the other about quality). Focusing on the first, other participants did also express concern about the *integrity* of data when it is owned and controlled by individuals. This ranged from poor accuracy and reliability (in the case of patient-generated health data), to more malicious things such as tampering with data. These are valid concerns.

Wolf et al. (2013) found that three out of four mobile apps designed to assess the malignancy of skin lesions, provided inaccurate results 30% of the time. In a more recent study, Sangers et al. found that similar apps powered by ‘Artificial Intelligence’ showed promise but are “far from perfect” (2022, para. 24)⁴⁸. There is, undoubtedly, great promise in the so-called mHealth market but skepticism remains around the diagnostic value of any data generated outside of the formal health system. This is summarised neatly by West et al., in the form of a scenario:

“If Rupert expected that the doctor would use his data to support a diagnosis, he would be disappointed because Rupert’s doctor cannot trust this data to support a diagnosis. Instead, the doctor may glance at Rupert’s data out of politeness, set it aside, and decide on which tests he [sic] would order to support a diagnosis, or refer Rupert to a cardiologist.” (West et al., 2017, para. 45).

There are a wide range of issues sitting underneath the acceptance or utility of patient-generated health data.

The counterargument to anticipated issues around patient-owned and controlled data, is that quality would be *improved*. The supposition is that, in providing individuals with agency around their own information, they will take greater care to look after it. The antithesis of this argument is that any centralised data relationship encourages users to become more passive around their data, since they feel they have no part in the process. This aligns with the discussion in section 2.2.3, where Malkin et al. noted that “people are learning that their data gets shared and repurposed, but they don’t necessarily condone it and they don’t feel like they have control over how much information is collected or how it is used” (2018, p.9).

⁴⁸86.9% sensitivity and 70.4% specificity in detection of premalignant and malignant skin lesions, to be exact.

7. Research Phase 4: Interviews

In support of this notion is the survey results we reviewed in chapter 4, where the accuracy of *existing* health information was criticised by respondents. Acknowledging once again that the survey was not demographically representative, this theme was relatively strong and suggests two key points which are pertinent here:

- There are data integrity and quality issues in the status quo centralised system, with one survey respondent noting that *“I would love the ability to actually correct the stuff that is written by the hospital. I find the amount of errors in the hospital notes I have seen utterly appalling [sic]. Things like the wrong year for surgery and wrong medication”*.
- Individuals have a desire to correct and update health information and, whilst there is a legislative right to do so, this right is under-utilised or confusing to individuals.

The third quote (*“The outcomes that we are seeking is higher quality data that captures more than an administrative process”*) focuses on the *usefulness* component of data quality, and expands our thinking about what data could be made available to share⁴⁹. For example, one of the core MOH datasets is the ‘National Minimum Data Set’ (NMDS) and this generally forms the basis for any health services research in Aotearoa New Zealand. Whilst it is a consistent and dependable data set, its genesis is in a funding and claiming context. On that basis, it has significant limitations when being used to make sense of population wellbeing.

A useful example of this is provided by Dawson (2020) in their assessment of the usefulness of government data sets in measuring maternal health equity. Specifically, the data point around whether or not a woman keeps her placenta is considered to be a key indicator of cultural wellbeing and hauora for wahine Māori. This data was actually collected by MOH until 2007 but, at the current time, is no longer available in any national data sets being relied upon for understanding of wellbeing. Dawson summarises this point by stating that “these data [government administrative data] are not being collected for what is important to citizens, but for administrative purposes and yet decisions that affect citizens are made from it” (2020, p.214). This is also reflected by an interview participant who stated that:

“There is hugely valuable information that sits in communities that will never ever sit in the administrative data that government collects and holds.” (Participant 342775a4).

Finally, the fourth quote suggests that there are workflow and business process issues which are contributing to a loss of data integrity. In the status quo scenario, centralisation creates an incentive to duplicate data, and this is aggravated by very poor or non-existent transfer mechanisms

⁴⁹We have also touched on this concept already in section 7.7.1.4.

7. Research Phase 4: Interviews

between clinical applications or data stores. In short, if you are lucky enough to be able to access some health data you require it will generally be made available as a *file* which will then often be imported into another system for usage. In the course of doing this, the data has been copied at least twice. This process introduces the risk that data can be modified at any of these steps⁵⁰ and, thus, a ‘drift’ from the source data occurs. If the shared copy of the data has to be updated in some way, there will generally be no write-back mechanism to the original source data and integrity is already severely compromised.

The quote above touching on this subject notes that having data in *one place only* solves this problem, but doesn’t hint at the technical complexity of making it happen. This approach can work in either a centralised or distributed paradigm, to be clear, and it is instructive to think about the innovation we have reviewed in section 2.2.4.1 which saw Estonia legislate that personal information can only be stored in a single place across government.

7.7.3.3. Empowerment

The ability of data to empower individuals rests on the assumptions that they want to access their data, they have the capability to utilise and interpret data and, in doing so, can have a less passive stake in their wellbeing. We have already discussed the issue of capability above, so now we turn to participant comments around the extent or likelihood of empowerment offered by a RDHIS:

“If people have agency to make the decisions for themselves, then that will: A, keep them more engaged and more involved in the process; B, it’ll hopefully mean that they make better decisions that are applicable for their own specific use case, rather than being subject to decisions that have been generalised across a whole population.” (Participant 4a626164).

Read literally this comment is actually about *decision-making*, which does not necessarily have to rest on a RDHIS. Any access to data, or even a more central and enhanced role in their care process, can result in the advantages that this participant touches on. My argument is that a RDHIS can *accelerate* this kind of increased engagement and agency, because the individual is such a critical part of the flow of data amongst all stakeholders⁵¹ and – critically – use and interpretation of any data can be made personally or culturally relevant.

This is supported by the survey findings we reviewed in chapter 4, and in general terms by the literature (Ministry of Health, 2016; Patel, Barker, & Siminerio, 2015; Woods et al., 2013).

⁵⁰This is especially a risk when Excel is utilised for any form of data capture or storage, which should be made illegal.

⁵¹In this case, stakeholders may include a team of health professionals, individual practitioners, whānau or community support groups – amongst many others.

7. Research Phase 4: Interviews

The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems have more clearly linked the concepts of data control with agency and wellbeing, and note that:

“To strengthen individual agency, governments and organizations must test and implement technologies and policies that let individuals create, curate, and control their online agency as associated with their identity. Data transactions should be moderated and case-by-case authorization decisions from the individual as to who can process what personal data for what purpose.” (The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, 2019, p.110).

This work does not refer to more architectural concepts such as centralisation or distribution, and yet it should be clear from all foregoing discussion that allowing individuals to “create, curate and control their online agency” is possible but problematic in a centralised paradigm. Other participants commented on the topic of empowerment in the following ways:

“Where you get a lot of ... agency is giving them the agency to choose when they share what with who.” (Participant 682d3989).

“The advantage [of decentralisation] is that it gives people agency.” (Participant 4a626164).

“But the more people engage and the more it is made part of their daily interactions, then far more people would see this is less about data sharing or consent giving action and more just about making decisions in their lives.” (Participant 682d3989).

So there is certainly sentiment amongst participants that control enhances decision-making, and increases *agency*. There is currently no literature available which might quantify a difference in perceived agency or empowerment between centralised and distributed paradigms. At present, there is certainly plenty of literature to support the general concept that access to data improves agency and flow-on effects around chronic disease management, for example.

Because decentralisation/distribution is such a new concept, we are only able to draw out potential functional differences between the paradigms and identify potential constraints of centralisation. However, we are not even in a position to gauge how well *centralised* data sharing and control performs, since even that has not yet been achieved⁵².

7.7.3.4. Incumbents

One of the interview guide questions asked participants to consider how incumbents in the health sector would react to implementation of a RDHIS. This is a broad question, since any such

⁵²Again, I should repeat that there are well-meaning efforts to increase individual access to centralised data but even this remains nascent and immature in Aotearoa New Zealand at the time of writing.

7. Research Phase 4: Interviews

change would affect a wide range of stakeholders – from IT system vendors, to population health researchers, to any health professional interacting with personal health data.

In this sense, ‘incumbent’ can be thought of as any person or entity currently operating within the *status quo* health system in Aotearoa New Zealand. This is a very important area to understand, since any change management around the philosophical underpinnings of our health ecosystem must include any incumbent and bring them along on the journey. A selection of participant comments is copied below:

“Your view on decentralisation falls on its face with [redacted name of popular Patient Management System]!” (Participant 6dea2c99).

“The incumbent systems and providers are going to need to preserve and retain their revenue streams.” (Participant 6dea2c99).

“It’s not in their [incumbent IT vendors] interest for it to happen, right? And so they will fight against it. In my mind, the argument is not yet strong enough. If the argument was so strong that this obviously leads to better outcomes for patients, then you have a stronger footing in which to stand.” (Participant 4a626164).

“There would be things that come from a different motivation, which are vested interests, be they professional or commercial or something else, in changing the power dynamics that exist around data.” (Participant 4a626164).

“That’s the shift we need in the mindsets of the vendors, is, within the bounds of privacy and security and giving people control, you’ve got to think about making that data accessible.” (Participant ca751e0d).

It is very understandable that, in the case of IT product vendors, a fundamental shift in the ecosystem, such as that represented by decentralisation/distribution, would not be welcomed. This does not necessarily have to be borne of malice or intransigence; vendors often have a long-term investment in a product, upon which the health of the entire business depends.

Influencing the health IT landscape is notoriously difficult (Gauld, 2004; Greenhalgh et al., 2010), although there is evidence to suggest that incentivisation can make a significant difference. For example, EHR uptake in US hospitals increased from 9% to 84% between 2008 and 2015, after the 2009 HITECH Act offered some \$20bn in direct incentives (Henry, Pylypchuk, Searcy, Patel, et al., 2016).

The above participant comments note that incumbent IT vendors, particularly, are not doing enough currently around making data accessible and, furthermore, would resist any attempts to change the underlying centralised operating model. The same issues were raised in 2016, when writing my Masters thesis. In discussing the influence of vendors there, I noted literature

7. Research Phase 4: Interviews

assertively identifying vendors as an interoperability *obstacle*. For example Terhune, Epstein, and Arnst (2009) wrote that vendors are *actively opposed* to interoperability (viewing it as a threat to business models); Koppel (2012) agreed and noted attempts by vendors to “balkanize” the market, and; Mandl and Kohane (2012) considered that vendors have consistently overstated the complexity of IT requirements in the health sector in order to maximise profit.

Whilst these citations are all US-based and more than ten years old, they do at least point to the degree of influence that IT vendors have over decision-making. Viewed in its most negative light (from a vendor perspective), the RDHIS has the potential to destroy their business model. Alternatively, vendors may already be looking towards a future based more on connecting data in place and *adding value* to that. With MOH signalling clearly in Aotearoa New Zealand that this is the desired future state, it will certainly be interesting to see how the vendor market responds⁵³.

In figure 7.5 the ‘incumbent’ subtheme is connected with two other themes. Its link to the ‘centralised hegemony’ subtheme reflects the difficulty of radical change in a complex and risk-averse sector; any change in health is slow, usually for good reason, but I intend to highlight here that the overall hegemony of centralisation is very strongly supported by the IT vendor landscape and the centralised models they have relied on and embedded.

Aggravating this is the notion that they have only been *lightly* regulated by the centre. One participant commented that:

“You’ve sort of been allowed to come to the point of view that you can control the data, even though I don’t think that’s an explicit setting. I just think that because nothing’s been said, people can just form their own view basically.” (Participant 6f5c5542).

So the vendors have understandably embedded philosophical and architectural models that support their profitability and, because of weak oversight from the centre, this has gone unchecked – potentially making it harder to unravel if that became necessary.

Finally, IT vendors are not the only incumbents. We have already discussed in section 7.7.1 that health professionals will understandably insist on retaining access to data they need to perform their functions. A shift in the ownership and control of data will necessarily affect health professionals too although, assuming we have addressed the general clinical risk issue, it is not clear which professional subgroups would actually object to patient control (if any).

I noted in 2016 that “knowledge as a business model” was an obstacle to the empowering of patients with their health information. This has a number of components:

⁵³One of the dominant Patient Management Systems in the primary care space has already move towards alignment with this indicative vision, by releasing an integration layer which has the ability to support a marketplace of third-party apps and integrations (MedTech Global, 2021).

7. Research Phase 4: Interviews

- The ‘fee for service’ business model of primary care
- The upstream funding model for primary and community care
- The view that empowering patients with information could be seen as an opportunity cost (Poor, 2016).

This is a much broader phenomenon than health, of course. Reeves et al. (2022) found anxiety and concern within the research community about the prospect of data sharing, specifically because it would affect the sole ownership researchers had over data they collected.

In the health sector there is an additional, more qualitative, aspect around professional autonomy to consider. There is literature identifying that health professionals – particularly doctors – are sensitive to any perceived impact on their autonomy (Bayless, 1996; Edwards, Kornacki, & Silversin, 2002; Swan & Newell, 1996) and, therefore, any rebalancing of power via data and knowledge must be viewed in this light.

7.7.4. Trust

Finally *trust* was prominent enough from the interview process to deserve its own theme, and rightly so. A large component of data sharing is based on trust (whether that’s on a bilateral and interpersonal level, or via some automated process). Trust cannot simply be generated by acting in a reliably trustworthy way; it is just as important that the individual has a *perception* of trustworthy behaviour (Belfrage et al., 2022) and this, of course, can be mediated by many other variables (such as social context or family/whānau perceptions, for example).

The interview process also threw up issues around the overlap, or conflation, of related themes such as social licence and consent. We have already discussed the interplay between social licence and consent in section 7.7.3 and so, in this section, I will focus on the higher order concept of trust.

7.7.4.1. Relationships are vital


“Relationships are vital” seems like an obvious statement to make and yet, for technologists, it can be easy to unconsciously disassociate the thing you are working on from its real world application – a humorous rendition of which is shown in figure 7.8.

In terms of data sharing and data sovereignty, it means that there is a risk we focus on what is *technically* possible without stepping back to understand the human-centred scenarios in which this data is generated. Participants commented on this in the following ways:

7. Research Phase 4: Interviews

Figure 7.8.: Automation (xkcd, n.d.)

"I SPEND A LOT OF TIME ON THIS TASK.
I SHOULD WRITE A PROGRAM AUTOMATING IT!"



"If we're incorporating data into most aspects of life, then being able to form and manage data relationships becomes a key part of what we do." (Participant 1655b226).

"We have to build it together. You start small and invite people to it, and build trust, and build relationships." (Participant 682d3989).

"We build relationships everyday, and they break everyday. Just think about it on a bigger scale. Behavior that is trustworthy and is maintained out of a genuine desire to benefit the other will be a healthy relationship that will endure over time." (Participant 682d3989).

"I think the biggest thing in building trust is relationships and intention." (Participant 1295cd16).

"People are harder than any piece of code you will ever write. The answers are simple, and there's often a reason we don't do them, and that's because they're hard. Relationships are hard." (Participant 1295cd16).

"It used to be very close in the early 20th century. People thought they were the government. Whereas now, I think it's further apart and more distant. There is less trust." (Participant df9eed6).

7. Research Phase 4: Interviews

“That master-servant type relationship, attitude needs to be deconstructed in order for real change for our people.” (Participant 1593596e).

In terms of the *trust* subtheme, there are a few key facets to pull out from the above selected comments:

- Trust is a product of healthy relationships
- Building healthy relationships at scale is difficult and complex because of the widely varying attitudes, opinions and world views represented in our society
- There is a sense of growing mistrust between the public and the government.

The notion that trust is part of a complex interaction of phenomena has already been shown in figure 7.7. In that representation, social licence is a product of trust. The additional layer identified from the above comments is that trust is a product of healthy relationships. We could argue that a healthy relationship (in a data sharing sense) is one where control has been delegated, individuals and groups are empowered, and use of data is verifiably in scope of any defined agreement. But perhaps we should consider this more broadly?

In 1973, Mark Granovetter published a ground-breaking study which attempted to define and quantify relationships. For example, he posited that the strength of a relationship can be defined as “a (probably linear) combination of the amount of time, the emotional intensity, the intimacy (mutual confiding), and the reciprocal services which characterize the tie” (Granovetter, 1973, p.1361). His conclusion that weak ties (more casual relationships with colleagues or acquaintances) actually have a stronger impact on social outcomes than strong ties (family members and partners)⁵⁴, has since been challenged in the context of the modern online era and the rise of social media (Krämer, Sauer, & Ellison, 2021).

It should also be acknowledged that there is a strong cultural component to this, as well, with te ao Māori prioritising concepts such as *whanaungatanga* and *manaakitanga* which have been overlooked by the centre when considering use of information and building data pathways (Bishop, 2016; Pool, 2016). This is reflected in some of the selected comments, but particularly one that notes a ‘master-servant’ relationship. Any description of a relationship with the government, or its agencies, in these terms indicates a very unhealthy relationship from which trust, and subsequently social licence, is unlikely to grow. This is also reflected in the comment about a general lessening of trust in government.

This is a large area of research which unfortunately we can only touch upon within the confines of this thesis. However, this brief discussion at least expands our world view to consider

⁵⁴This conclusion rests on the idea that weak ties can provide access to information that is not found within the individuals’ network of strong ties.

7. Research Phase 4: Interviews

all the ties and relationships that people might hold, and how they are perceived and utilised by each individual. This certainly does not make the topic less complex for data sharing, but enhances our understanding of it.

Finally it should be noted that we have only covered the population relationship with government at large so far. Clearly, society is comprised of a much richer and more complex mesh of trust relationships than only this one. The trust in government component has been highlighted because currently the flow of data is, generally speaking, under the control of government agencies. But, certainly, there are a wealth of other trust relationships at play. One particular comment is instructive in terms of how a future data ecosystem might unfold:

“Māori people are going to trust iwi much more than they will trust the government when collecting data.” (Participant edc46a04).

Certainly, there are many iwi and not all iwi have close relationships. But, if we can accept that a general level of trust is higher between Māori and iwi than Māori and government, then it helps to understand the data sharing landscape within Aotearoa New Zealand⁵⁵.

7.7.4.2. Decentralisation's impact on trust

Following on directly from the last section, some interview participants considered that a RDHIS could be a vehicle to generate more trust. This rests, perhaps, on the notion that permitting individual ownership and control of data is a concession of sorts, and one which has the potential to put government agencies at a relative disadvantage⁵⁶. We have already discussed the issue of government's need to access data in section 7.7.1. When we are considering trust, it goes some way to describing the hegemony of centralisation that any move to unravel that could be perceived as a trust-building exercise. Participants commented on this subtheme in the following ways:

“It [decentralisation] unlocks a willingness to take a risk and trust, because it doesn't take one direction.” (Participant 1655b226).

“In order for that trust to be given, transparency has to be there. How better to give transparency than to give control to the hands of people.” (Participant 682d3989).

“Trust is one of those things, you have to give it away to get it.” (Participant 682d3989).

⁵⁵However we should also recall the WOCA data request incident (Neilson, 2021), and be extremely cautious about discussing the opinion of Māori as a homogeneous group.

⁵⁶Relative to the status quo where the centre has almost exclusive access to all centralised data.

7. Research Phase 4: Interviews

“I think the idea that it’s decentralised has a benefit in terms of helping us address those people that don’t fundamentally trust their data being held centrally.” (Participant 6f5c5542).

The above comments, overall, describe a situation where it is reciprocity (“it doesn’t take one direction”) and the ceding of total control over centralised data (“you have to give it away to get it”) that increases trust. In fact this even links us back to the quote from Granovetter (1973) who considered that reciprocity was a supporting factor in strong relationships, and perhaps it is a perceived *lack* of reciprocity that aggravates some of the other variables contributing to mistrust. We saw, for example, in section 7.7.1 several participant comments bemoaning lack of access to centralised data in general.

There is another perspective on the impact of decentralisation/distribution on trust, however. One participant noted that:

“One of the core weaknesses of a decentralised approach is ... that it relies heavily on trust. So basically, if you set the rules, you rely on the other party you interact with to also follow those same rules ... there’s a lot of trust that you’re not going to do the wrong thing with that data, and that comes down to the fact that data is non-rivalrous.” (Participant 4a626164).

This is a very interesting comment since it raises the prospect that decentralisation/distribution, improperly utilised, can negatively impact trust. In this scenario, the recipient of shared data uses that data maliciously and the original owner has no way to prevent that.

It is certainly true that, once you have access to data, there is nothing that would stop you from copying it or emailing it or posting it on social media. Similarly, there is nothing that can *stop* a person from driving on the wrong side of the road, or stealing cigarettes from a convenience store. Trust is a slippery notion, but is clearly one that the effective functioning of society relies upon in lots of soft and imperceptible ways.

This may well be a false analogy, however. It is literally and physically impossible to stop people from performing certain actions. But it *is* possible to stop people having access to data – a situation also known as the *status quo*. Following the line of reasoning in the above quote, we may want to end our interest in a RDHIS because there is a possibility that a data sharing relationship will be abused. Or, alternatively, we go back to the primacy of relationships – in all their shapes and forms – and accept that, since it reflects the vagaries of being human, we will empower people to develop high-trust relationships and determine their own bilateral or multilateral agreements as they see fit. Empowering people does not come without risks, to be sure, and, at the very least, this comment raises the stakes around capability and education that we touched upon in section 7.7.2.

7.7.4.3. Social licence as a product of trust

We have already discussed the relationship between relationships, trust, consent and social licence in the preceding sections. Here, I would like to highlight participant comments that drew out some themes around social licence specifically. There are two components to this. Firstly, levels of social licence will vary significantly across all groups and subgroups within society and it is not a tick-box task. Secondly, that social licence can never be assumed even under the most ostensibly positive use of data.

Selected participant comments on the first component are reproduced below:

“You have to be genuinely open to the notion that maybe you can’t do the thing you wanted to do. And that’s something that decision makers don’t always honestly engage with because they view the consultation part as maybe being more of a box-ticking exercise. They just say, “We held a hui, everybody heard about it, we can go ahead”.” (Participant 4a626164).

“With social licence, it obviously changes over time and it does shift as well ... you could argue that social licence is not ubiquitous. It comes down to generations – Gen X, Gen Y or Baby Boomer. You can’t be arbitrary about it, you have to be granular about it, I suspect.” (Participant ca751e0d).

What we see in the above comments is both a recognition that societal groups are rarely, or even never, homogeneous. We can certainly interpret trends around social licence based on aggregated data but, in terms of implementing data sharing or data use initiatives, I am arguing that the only truly accurate gauge of social licence is the proactive and informed consent of individuals. It must stand to reason that anything less than this is simply making an assumption about the social licence conferred on the organisation concerned by all members of the group. This happens to be the *status quo* with most data use initiatives today.

It is certainly true that in the centralised paradigm we currently witness in health, seeking individual consent for data use *ex ante* is an impossible task, and I do not underestimate the effort required. Nevertheless, because of this the system is used to almost crossing its fingers and hoping that potential objectors do not notice the data use has occurred. As one participant has noted:

“[Under decentralisation] there wouldn’t be the same uncertainty that perhaps there is now about, you know, are we able to do this legally, do we have social licence to do this, how do we weigh up whether or not we should be doing this with data that we just happen to have.” (Participant 6f5c5542).

7. Research Phase 4: Interviews

So there is an acknowledgement that the social licence component is cumbersome and difficult to implement in the current system, and there simply isn't the resource or capacity to undertake that process to meet the naive idyll I have offered. This phenomenon is also reflected in the literature where the regular use of linked administrative data sets, even without specific consent, has become normalised such that it "enables maintenance of the social licence required for these resources to be used without specific individual consent" (Dawson, 2020, p.211). That is, use of data without social licence has become normalised and, at best, agencies will feel compelled to attempt to retrospectively attain that social licence.

This then further relates to the second comment, which expands on this and notes that even meeting representatives from key groups is far from sufficient – the exhortation is to be "more granular". This makes intuitive sense. Levels of trust, and affiliated social licence, are no more likely to be homogeneous across any social grouping than ones politics or favourite character from *The Sopranos*. The comment identifies age groups as one potential boundary but, even *within* Gen X alone, there will probably be as much variation in opinion as any other group.

Again, I am not suggesting that identifying themes from aggregated research is pointless. It is certainly useful to know, in general, whether or not older people trust government agencies. However, the point of this conversation is that *individuals* should be indicating social licence themselves via the giving or withholding of consent. When it comes to use of individual data, no other proxy or rounding mechanism will do.

The second component of this subtheme is described by a participant as follows:

"Even though there is a benefit to citizens from using integrated data at an individual level, that doesn't make it right and it doesn't mean they're going to like the idea. There are so many assumptions around – well, I think it's for their good. But is it my call to make?" (Participant 342775a4).

This is the only comment made on this particular point, but I think it's a very meaningful one. It speaks to a mindset where someone genuinely has the best intentions; they may be a socially conscious and politically aware data analyst who has an idea to combine data to produce something they feel will make a positive difference. Much of the discussion in this space before has been on malicious or unauthorised uses of data, but this is quite a different perspective. When operating from this mindset it might be tempting to overlook ethical issues such as social licence, or even whether the data use accords with its purpose under the Privacy Act. The message must be that even the most ostensibly positive use of data should be rigorously questioned from an ethical standpoint. Yes we can technically do it, but *should* we?

Once again, this is an issue that the RDHIS entirely solves via the virtuous circle depicted in figure 7.7. Trust begets social licence, which begets a defined purpose under the Privacy Act,

7. Research Phase 4: Interviews

Table 7.9.: Key findings from Research Phase 4

Topic	Description
Access to data	The centre has a monopoly on access to administrative data, and groups are excluded from generating their own understanding of their data.
	Government and health providers need access to data.
	Distribution has the potential to offer <i>more</i> data, when levels of trust and social licence are high.
Capability	Concern around capability of government and individuals if moving from the <i>status quo</i> .
Social licence	Accepted as crucial, yet sits very uncomfortably within a centralised paradigm.
Control	Will empower users/groups and threaten incumbents and vendors.
Trust	Distribution clarifies links between trust and social licence in a way that centralisation cannot.

and then consent is provided. It should be very clear under what terms the data is shared and how it can be used – the removal of ambiguity drastically minimises the ethical permutations of a proposed use of data.

7.8. Summary

The interview participants generated many valuable insights, and a great deal of ground has been covered in this chapter. The aim has been to understand more about the real world context in which the prototype might be applied and we have indeed identified practical obstacles to any implementation of a distributed information ecosystem in the New Zealand health sector, much of which is in fact generalisable to other sectors. I have summarised the key findings from the preceding section into table 7.9.

What this tells us is that there are very real issues with access to data. The status quo provides access (albeit poorly) to government and health providers, but results in an equity conundrum which means groups are disadvantaged (Waitangi Tribunal, 2019) and the links between trust, consent and social licence are needlessly ambiguous. I discussed in section 7.7.3.1 how these things become conflated, and demonstrated that something like a RDHIS clarifies them and empowers the individuals and groups who are defining those data-sharing relationships. I will

7. Research Phase 4: Interviews

draw all these threads together in the next chapter.

8. Discussion

This thesis aims to develop understanding about how data sovereignty can be achieved. I have recognised that, whilst the topic as it relates to Māori is currently being actively debated within the Aotearoa New Zealand public sector, very little attention is being paid to how it can optimally be achieved on a technical level. This is part of a global pattern where, I argue, two key factors are in play:

- The conceptual power of the centralised hegemony, which still dominates and constrains any discussion around radical changes in data ecosystems
- A lack of direct incentives for agencies and organisations to do things differently.

An intended output from this research is to demonstrate that truly distributed approaches are rapidly maturing, and it is eminently feasible to consider building them today. In so doing, I would like to draw attention to the fact that *another way is possible*. Given the two points above, a radically distributed health information system is most likely to be initiated outside of the state apparatus. It may be that Māori groups explore the possibilities for technology such as Holochain to provide self-determination in a way that is meaningful to them, and this further influences public sector trends. It seems unlikely that a private sector vendor would engage with this concept because of the inherent lack of ability to monetise it. It is difficult therefore to see what kind of a future there may be for radical distribution of data, but I maintain that it represents a critical vehicle for trust, social licence and civic engagement. The Aotearoa New Zealand government has realised the importance of social licence in its Data Protection and Use Policy (Social Wellbeing Agency, 2022), and I have argued that allowing people to own and control their own data is the most effective way to build trust and improve social licence.

In this chapter I will draw together all the different research phases, and return to the research questions introduced back in section §1.3. I will provide an overview of what I have learnt throughout this journey, and highlight what I consider to be contributions to the knowledge base. Finally I will provide an overview of possible directions for future research.

8.1. Research questions

8.1.1. How do people feel about the ability to have absolute control over their personal data?

In chapter 2 we saw that, for the general public across a range of jurisdictions, there is a fatalism around a loss of control regarding personal data. Thanks to the work of the Data Futures Partnership (2017a) and the Social Wellbeing Agency (2022) we understand that this is not confined to social media, but applies equally to public sector data capture. In the Aotearoa New Zealand setting, a growing data sovereignty movement has raised these issues in parallel in terms of how they relate to Māori and the consequences of colonisation. In this respect, people are starting to talk more about how they would like to feel a recalibration of the power imbalance seen in government collection and use of data.

Primary research conducted for this thesis certainly had limitations, as we have reviewed in chapter 4. The survey participants were not demographically representative, but nevertheless provided interesting insights, with the most pertinent conclusions being:

- People want to exert control over their health data
- People want to know who has accessed their health data
- People have very high levels of trust, in terms of data sharing, with their GP
- People have low levels of trust in government agencies.

This last point speaks to a social licence issue. In section 7.7.3 I proposed a model whereby a virtuous circle of data sharing is initiated by individuals having trust in government agencies. With the control and insight offered by a distributed health information system, this could result in the win-win of government having access to a range of high quality data which is willingly shared by citizens who have trust in its use. Given the trust issue noted above, this will certainly take some work but, in my view, is the optimal way to move forward. The MOH *Hira* project aims to provide some transparency and oversight of health data, and its designers are certainly to be commended for that. However, as already noted, where a model is based on a centralised architecture there is likely to be not only a risk of re-centralisation, but a *tendency* towards re-centralisation.

This research question was also addressed during the interview research phase (section §7.7). Some additional themes came to light there, with respondents discussing how the current flow of data into government is unidirectional and it is very hard for individuals, groups, or NGOs to access any health information at all. There are definitely legislative considerations at play

8. Discussion

here (Neilson, 2021), but participants expressed concern about the inability of groups to utilise health information to draw their own meaning, provide insights that are inclusive of ethnicities and sub-groups and permit groups to make inferences that are meaningful to them. I categorised this as ‘analytical sovereignty’, and maintain that this is a vital component of social licence. But there should be a pathway for anyone exercising analytical sovereignty to raise the profile of their findings in a meaningful way, lest they simply be ignored in favour of the *status quo*.

Another competing issue here is around capability. Many interview participants expressed concern about the prospect of an individual user having to negotiate access rights to every different component of their health data. I believe that the aim should be to offer this capability for those who want to exercise it, however I had to remind those participants that this is merely a prototype proving a technical concept. The detail around how an app might work, and what options are available, are out of scope. But, certainly, capability is an issue that will have to be addressed but there are established User-Centred Design principles which can cater for this.

What is also clear from the available literature (although it should be obvious) is that there is no homogeneous attitude we can attribute to a group. Attitudes change across every conceivable variable. By locating this idea within the data sovereignty debate, however, we raise the profile that it is marginalised groups who, on average, will be motivated to pursue some kind of data sovereignty and will benefit most from it. Even as a particularly zealous privacy advocate, I would find it difficult to argue that personal data sovereignty is in the realm of a *human right*. However, it is the consequences of not having access to data – and not holding analytical sovereignty – that are marginalising groups, contributing to negative outcomes, and thereby potentially impacting their rights.

A high-profile example of this can be found in the WAI2575 Waitangi Tribunal report, which found a breach of Treaty principles in the way that, amongst other things, health data was not used optimally to understand Māori health outcomes and to elucidate the causes of inequality (Waitangi Tribunal, 2019). Where data is to be considered a product of cultural knowledge, or a *part* of cultural knowledge, it is conceivable that Treaty claims could be extended to question why it is not being devolved to iwi or hapū, and why there is no mechanism to effect the analytical sovereignty that WAI2575 implies is needed.

The literature overall points to a desire for greater control over individuals’ personal data. However – and I recognise that this may seem self-serving – the absence of a strong movement clearly saying “we want control and ownership of data”, even amongst groups such as Te Mana Raraunga, is more likely to be a product of the centralised hegemony than anything else. This cannot be proven, of course, but it stands to reason that something which is at present so nascent and limited to mostly fringe application is not going to yet be influencing public opinion in the way that other technological advances have. This is furthered, I contend, by a verifiable absence

8. Discussion

of incentives for government agencies and organisations to proactively explore this complex arena. Most jurisdictions do not have a mobilised, engaged and expert group of people with “skin in the game” to raise the profile that such a thing is possible. It is very understandable that, left to open source interest groups casually developing pet projects on the internet, a distributed personal health data movement will not gain a foothold in, say, the UK or Japan. Aotearoa New Zealand, on the other hand, has the rare combination of both an active thought-leadership group and constitutional tools that can be brought to bear.

8.1.2. How can technology support the distribution of personal health data and a move toward data sovereignty?

In this thesis I have positioned personal data sovereignty as the overarching goal. This has been discussed in literature by indigenous researchers, but has a very low profile elsewhere. I have argued that the ‘distributed’ data model (Baran, 1964) offers the most possibility in enabling data sovereignty. Until now, the question has been *how do we get there?*

I am not an indigenous researcher and have no right to contribute to the indigenous discourse, but I have been inspired by the thoughts and concepts behind indigenous data sovereignty. At the same time, my technical interests have made me aware that there are maturing technologies which offer a way to implement *practical* solutions for data sovereignty. The key research output of this thesis is to prove that a distributed app can be used in a healthcare setting. In chapter 6 I describe the development of the V2 prototype and its evaluation process. The user requirements that had been elicited empirically were all met. At a minimum, therefore, I have proven the basic feasibility of a distributed app which would represent an enabler for data sovereignty. This prototype app is available to anyone to download, install and evaluate themselves¹; the full codebase is also available for inspection².

An important question that arose during the interview process is whether or not a technical solution is even required; do we actually need a tangible solution for data sovereignty? We saw that relationships are vital and, connected with this, people need to feel trust in how their data is used. That trust appears to be low (Data Futures Partnership, 2017a; Social Investment Agency, 2018). We also saw an initial focus from Māori data sovereignty researchers on governance. I certainly do not underestimate the importance of relationships and governance and ‘soft’ variables – such as trust and social licence (after all, these comprised a large proportion of the discussion in chapter 7) – but other literature has noted the risk of a tendency toward re-centralisation, if it is technically possible (Gottsegen, 2021; Walch, 2019). This implies that some formal boundaries should be in place; a hard solution is necessary, and I have demonstrated

¹<https://gitlab.com/alexpoor/radhis/-/releases/v0.0.2>.

²<https://gitlab.com/alexpoor/radhis>.

how such a solution might work.

8.1.3. How could a distributed health information system be implemented in Aotearoa New Zealand?

Building upon the previous two research questions, we are now at a point where we can reasonably ask – *what next?* It is one thing to understand that groups – even, potentially, large groups – desire data sovereignty, analytical sovereignty and being able to profess a narrative which is appropriately grounded in their own lived experience and backed by their own data. It is another thing to demonstrate that this is now technically feasible. It is *yet another thing entirely* to understand what is required to make it a reality! This was the aim of the interview research phase, discussed in chapter 7, where the outputs could be summarised as a roadmap. In order to do so, I will first reiterate the key *practical* issues that arose from that research phase, before drawing them together into a whole.

8.1.3.1. Identity

This refers to the practical requirement to understand who people and organisations *are*. In the demonstration video (<https://vimeo.com/680251596#t=949s>) this is glossed over by having users pretend they are health professionals. In the real world, it is vitally important for my GP to know that I am really Alex and *vice versa*.

The issue of identity is universally important in a healthcare system, and distribution of data does not make this less so. In Aotearoa New Zealand there is a very high-quality health identifier – the National Health Index (NHI) – which is used in all health system interactions³. The *Hira* project has the stated aim of opening access to core MOH services, such as NHI, for use via API, and rapid innovations during the Coronavirus pandemic have proven that individual identity can be verified and connected with a health record⁴. There is no reason why a RDHIS app could not leverage the same API service and verify a user's identity to link it with the correct NHI.

The provider side is more complex although, again, MOH do have some services in place. For example, an official health facilities list is available for download⁵, and contains 12973 officially recognised health facilities (or providers) as at 9 May 2022. It is conceivable that a RDHIS app could access this list to present an index of providers to a user, and ask them who they want to share data with. For other data sharing use cases – for example, to share with a friend or

³Although most people would not know their NHI, it is a vital data point for finding the correct patient information and healthcare administration.

⁴The MyHealth service permits account creation either via RealMe or using official identification.

⁵<https://www.health.govt.nz/nz-health-statistics/data-references/code-tables/common-code-tables/facility-code-table>.

8. Discussion

family member – there would be no need to access a centralised service. Some workflow could be developed allowing Near Field Communication (NFC) or QR code generation, which would connect two users and allow them to share data.

Without getting too much into the solution space, identity needs to be highlighted as a critical area that needs further consideration for implementation of a distributed system. For any interactions which cross into the formal public sector space, access to centralised reference data is unavoidable.

8.1.3.2. Access and equity

Before embarking on any digital implementation programme such as proposed here, it is vital to assess the potential impact on population groups. I have repeatedly referenced the perceived benefits of a distributed system as accruing to marginalised groups, who would have the opportunity to better reflect their lived experience whilst basing it on data. It is also important to acknowledge that those same marginalised groups may in fact be the *least* equipped to actually engage with analytical sovereignty and in some cases may not have the access to technology to engage with a RDHIS at all. There is a risk, therefore, of making things *worse*.

This point simply serves as a reminder not to assume anything about access and equity, and to always be conscious of unintended harm. MOH will progress with their *Hira* programme, which will also take these factors into account.

8.1.3.3. Trust and social licence

Trust was identified as a separate theme in chapter 7 and, in summary, there was a prevailing sense of low trust in government use of data and a desire for greater effort to be expended in its corollary, social licence. This may be considered as something of a ‘holy grail’ for governments who wish to have confidence and transparency in their use of centralised population data. I argue in this thesis that trust and social licence form part of a virtuous circle, which must be initiated by first *ceding* power. This introduces the risk that government will have access to less data than at present; this was a concern echoed by most interview participants. There is also an interesting tension here, when we consider the WAI2575 recommendations were about *ineffective* use of data to understand equity which has resulted in poor health outcomes for Māori (Waitangi Tribunal, 2019). Taken in isolation, this might embolden the centre to *strengthen* centralised control of data. Conversely it may also be seen as an inflection point where, as a system, we want to redesign the data landscape from the ground up to move towards trust, social licence and better social outcomes via distribution of data.

The proposed ceding of power as part of this process will be very difficult for the centre to

8. Discussion

deal with. Similarly, reduced access to data will have operational and strategic implications for MOH, Health NZ and the government. It will be vital to ensure that as much work as possible has been undertaken to minimise the impact of this.

8.1.3.4. Legislation

The Privacy Act 2020, by virtue of its flexible consent-based model, seems well placed to cater for distributed data flow across the public sector. As we have discussed, the Act offers no obstacles to implementation of a RDHIS. The work to be done on trust and social licence will substantively alter the landscape of ‘Purpose’ under the Act, by allowing individuals and groups to define their own Purposes when sharing data – and they would also have the ability to check compliance with the defined Purpose, well before a privacy incident has surfaced.

However, there are numerous issues with domain-specific legislation in the health sector. For example §45 Medicines Act 1981 mandates that prescribing information is kept securely at “his [sic] place of business” and this information can be inspected and copied by “any officer of the Ministry of Health”. Similarly, the Health (Retention of Health Information) Regulations 1996 stipulates that all providers must retain health information for ten years.

This is simply to highlight the fact that there are many legislative instruments that affect health sector operations, and the substance of each will need to be reviewed and amended to accommodate any move away from centralised data. This need not be a legislative paradigm shift. For example, §6(1) Health (Retention of Health Information) Regulations 1996 notes that “every provider *that holds health information* [emphasis added] shall retain that health information for the minimum retention period”. It would be possible to interpret this section differently if health information were not *stored* with providers by default, but kept only with the patient. Conversely, the same Regulation (§7(2)) extends to anyone holding information “for the time being”, which could be read into the scenario where data is shared. There is certainly work to do in this space.

8.1.3.5. Critical access to data

As noted above, a strong theme during the interview research phase was anxiety about government retaining access to data under a RDHIS.

The literature, as well as the survey discussed in chapter 4, strongly indicates a high level of trust in the primary care team. As the routine interface between individuals and the health system this makes sense. Going further, the ongoing review of the Aotearoa New Zealand health system aspires to keep people well, and employ interdisciplinary models of care which are meaningful to the locality in which they live (Health and Disability System Review, 2020). This strongly

8. Discussion

implies an enhanced role for primary and community care moving forward.

It therefore makes sense that some default is enabled in a distributed system, where the primary/community care team do have access to a patient's data. Patients could proactively opt out if they wish. Secondary uses of data – most commonly for claiming and funding – could also be enabled by default from the primary care provider site. The important difference from the status quo is that all secondary uses of data would be visible and auditable to the patient. Put another way, there would be a 'baseline' of critical data required for the health system to function which could be activated by default. This would also encompass genuine emergency use cases, for example by paramedics.

Critics will argue at this point that things do not seem very distributed at all. We have data being sent to the primary care team and beyond. However there are two important rebuttals to make. Firstly, all secondary use of data is visible and auditable. Secondly, distribution represents an important ontological shift in the balance of power. I have theorised that this is critical in building trust and social licence, which continues to be something that all of government is concerned with (Social Wellbeing Agency, 2022).

8.1.3.6. Analytical communities

Finally, the promise of 'analytical sovereignty' via distribution of data requires that we imagine a different operating model. Currently all data flows to the centre, where it is analysed and used for policy development. The Waitangi Tribunal (2019) found that use of data was not being done effectively, however, and contributed to negative health outcomes for Māori. The potential advantages of distributing data include allowing it to be held and utilised by groups or communities. These do not need to be prescribed; they can be anything that people want them to be. The RDHIS should include a facility for people to identify as part of a group or community; they would be able to define a data sharing arrangement with that group.

In an ideal world, the group might be a collective of individuals who are traditionally marginalised – for example, trans or disabled people. The group would collate data on behalf of its members, and be able to exert its analytical sovereignty by conducting analysis on that data which is meaningful to them, and reflects their lived experience. Certainly, not all groups would have this capability and it may be that the groups themselves become part of larger groups, who could be NGOs with more analytical capability and whose interests align or overlap with the underlying group.

The point here is that the outputs from use of this data need to be fed back to the centre. MOH/Health NZ retain a core policy advice role to government, and they need to consider these analyses. I venture to suggest that this will be a difficult shift for them; the centre certainly has access to highly-skilled people. But WAI2575 taught us that this is apparently not

8. Discussion

enough. Therefore, implementation of a distributed information system will need to coincide with a changed operating model, whereby groups and subgroups are self-organising and self-determined, but there is a formal pathway for their findings and analyses to be considered by government. Note that, initially, this will occur in an environment where the government is likely to have reduced access to data. They should therefore be highly motivated to understand and consider this new source of intelligence.

8.2. Implementation model

We can now piece the foregoing discussion together to draft an implementation model for a RDHIS. Please note that this is intended to be a set of primarily practical considerations, which can be used as a reference for others wishing to understand what is required at a high level.

The implementation model is shown in figure 8.1. The model reflects the key points from prior subsections in this chapter, but groups them together to be used as a reference guide for implementation of distributed data systems in a public sector. The language is purposefully generic so as to be generalisable to other jurisdictions if necessary.

Utilising this model would place a public sector system in a good position to maximise the benefits of moving away from centralisation. However, this obviously does not tell the whole story. Something that's worth drawing out is the operating model for analytical communities.

These are self-determining groups (and potentially subgroups) that are engaged in utilising data from that collective; they exert analytical sovereignty on behalf of their members. These groups are important, critical even, in any distributed data ecosystem. Initially, however, they may be *more* important since the state apparatus will not have the access to data it is used to. The operating model within a given sector will therefore need to reflect this; by simultaneously accessing data from those individuals providing bilateral access, and also by formally utilising analyses from analytical communities. A representation of this is shown in figure 8.2.

8.3. Significance of research


This thesis makes an original contribution to knowledge by extending the literature on data management models, building a prototype app which proves the feasibility of distributed data in healthcare applications and also contributes a model. The research overall should prove a useful starting point for anyone wishing to investigate applied uses of distributed data in the public sector. I will now briefly review what I believe makes this research significant.

8. Discussion

Figure 8.1.: RDHIS implementation model



Figure 8.2.: Analytical communities



8.3.1. Contribution to the knowledge base

This is the first piece of research to investigate uses of distributed data in the Aotearoa New Zealand health sector. The findings are sufficiently generalisable to apply to other public sector domains, and potentially remain valid for application across different jurisdictions with a broadly similar political base. If the US federal government, for example, decided that it wishes to explore distributed data then – although the health system is fundamentally different in almost every conceivable way – all the steps shown in figure 8.1 remain valid.

While there is a large research base around blockchain and decentralisation (including many studies within health systems), to the best of my knowledge none have yet looked at the distributed model (Baran, 1964) or the maturing technologies which are leveraging peer-to-peer concepts. These technologies are relatively new, but are developing rapidly. This work also adds to a very small extant research base of projects utilising Holochain in differing ways.

The prototype app is certainly limited in its capabilities, and there is a much broader programme of work that must follow, but it also represents a contribution to the knowledge base in its own right. The app itself, and all the code I have written, is freely available at <https://gitlab.com/alexpoor/radhis> (under the Cryptographic Autonomy Licence 1.0). Furthermore, it can be installed and utilised by anyone for testing or evaluation.

While I have tended to focus on the prototype app as the ‘instantiation’ artifact, we must not forget the work of March and Smith (1995). As shown in table 3.1, they in fact define four separate types of artifact which DSR can produce. I consider that this research has produced three of these artifact types:

1. **Construct.** The reader may think I am being too self-congratulatory to declare the literature review as an artifact. But the DSR literature is in agreement that this is exactly what it is – when done purposefully. The construct is derived from existing knowledge and helps us to understand the problem. In chapter 2 I explored the semantic confusion around terms, how they are being misused, and arrived at what I hope is some definitional clarity between centralisation, decentralisation and distribution. This taxonomy itself dates back to 1964, so it is certainly not new. But we are now seeing an increasing trend of ambiguous concepts, which all overlap. I am confident in saying that overlaying the current state of selected technologies with Baran’s taxonomy is a very important research artifact, without which the thesis would be pointless and which I hope assists further research in this space.
2. **Model.** Models help us to understand solutions and problems and connect them both by allowing us to explore the changes they would effect in the real world (Hevner et al., 2004). I propose that the implementation model shown in figure 8.1 represents such a model. Both the problem and the solution space is defined and informed by all stages

8. Discussion

of the research. The implementation model is simple to comprehend, yet is based in empirical research conducted purposefully.

3. **Instantiation.** Finally the prototype app represents an instantiation, whereby both the construct and the model are given life in a new product and feasibility is demonstrated.

8.3.2. A call to action

For me personally, none of the above would be very meaningful if it could not have a real world impact of some kind. I have repeatedly noted the dominance of the centralised paradigm, and how I believe it is limiting our thinking around what is possible. Having worked with data in the public sector, I feel strongly that things could be done much better. WAI2575 demonstrated that the large amounts of data the centre are already getting are not being used effectively, and has directly contributed to unequal health outcomes for Māori. Distribution of data can solve this problem by giving Māori sovereignty over their data, and respecting their ability to derive culturally appropriate meaning from it. I contend that this cannot be functional or sustainable in a centralised paradigm.

Similarly, there are other marginalised groups – and non-marginalised groups who are simply interested in having authority over their personal data – who can benefit from a similar approach. The case of vTaiwan (discussed in section 2.2.4.4) has shown us that, given the ability and capacity to do so, citizen groups are motivated and capable to build new artifacts of their own and contribute to sense-making and decision-making. The Aotearoa New Zealand government's commitment to 'open data'⁶ is charming, but ineffectual for any kind of analytical sovereignty. For example, mortality data sets are available but none more recent than 2015 (seven years old at the time of writing) and they provide only the ability to, at the finest level of detail, identify counts of the main cause of death in a DHB region. It is easy to imagine subgroups becoming lost in the 'noise'.

I would therefore like this thesis to also be the start of a discussion within Aotearoa New Zealand, about how we reclaim our position as a data and digital thought leader and encourage the kind of "hacktivism" seen in Taiwan to purposefully improve social outcomes. I acknowledge this sounds idealistic, but we know now that data can be effectively distributed and there are jurisdictions who are inspiring us to think differently about civic engagement and transparency.

8.3.3. Possibilities for further research

This thesis poses more questions than it answers, perhaps. Having overcome the hurdle of demonstrating that health data can be distributed (on a technical level), a slew of follow-up

⁶NZ signed up to the International Open Data Charter in 2017 (digital.govt.nz et al., 2020).

8. Discussion

questions come into play. Some of these have been encompassed in the implementation model (figure 8.1), but only at a high level. I have kept a careful track of these follow-up questions, as they have occurred to me or been suggested by the research participants, and I attempt to cohere these into a logical sequence below.

Holochain

- Investigate making full use of a Holochain DHT by storing encrypted health data *on* the DHT itself. This has a lot of benefits in terms of data availability and resilience, but requires a different focus on cryptography and Public Key Infrastructure (PKI) workflow.
- Test feasibility of ‘querying’ aggregated data from a range of distributed sources. This is required to make analytical communities practical, in essence by extending the ‘X-Tee’ approach down to individual users and devices.
- The security of the Holochain app should be tested rigorously, to understand more about potential attack vectors for personal health data. This was raised as an issue during the survey research phase, and it deserves investigation in the context of a distributed app specifically.

Data sovereignty

- Conduct focused research into the clinician’s perspective on distributed data sovereignty. What practical implications are there for a clinical workflow in different settings?
- Investigate the impact of distributed data sovereignty on vendors currently operating in the health marketplace. There is an interoperability component here, which will help inform time-to-value.
- Further work is needed on decentralised/distributed identity management. Some technologies (SSB) provide an informal solution, but this is an area that deserves further exploration.
- Qualitative research conducted with communities and groups is needed to help understand the potential of analytical communities. This research should attempt to understand what kinds of groups will self-organise, the issues that are most important to them, and their readiness to engage with distributed data.
- Similarly, a comparison of analyses by different groups – to demonstrate variation in understanding, process, and findings – would be fascinating and would provide further evidence of the need for analytical communities.

8. Discussion

- Unanticipated negative impacts of distribution on vulnerable people should be considered (for example, people in abusive relationships or experiencing family violence).
- A sense of the ‘limits to distribution’ would be very interesting. For example, if we accept that it could be done in health (not necessarily easily, but it is feasible) then what other parts of the state apparatus would be considered off limits?

Legislation

- A detailed review of domain-specific legislation in the health sector (and other sectors) is needed, to clarify exactly where legislative obstacles to data sovereignty exist.

Trust and social licence

- Further understanding of public trust in government agencies around use of data would be welcome. Whilst literature does exist on this in the Aotearoa New Zealand context, no such work has yet been positioned in the context of data sovereignty and based on the knowledge that there is a technical route to achieve it. Reframing the conversation with this in mind would represent a very important contribution to the knowledge base – more so where the sample population is demographically representative.

8.3.4. Conclusion

As (I presume) with many PhDs, this thesis has been quite a journey. I have learned a lot – not only in the fascinating literature I have read and attempted to represent, but in the hard technical skills I have developed in building the prototype artifact. This thesis encompasses everything I am passionate about – technology, equity and social justice – and I had a lot of fun putting it together.

Having said that, I strongly believe that access to data is something that has the potential to right historic wrongs and uplift the whole of society in Aotearoa New Zealand. It will certainly be difficult; the ceding of power always is. But it is necessary and something that is actually easy to do, once we have reconciled ourselves to it. Giving people access to knowledge and data that allows them to represent themselves in the analyses and findings they make – but also to turn that into tangible outcomes to make their lives better – is the awesome power that data has.

New technologies like Holochain offer the potential to reshape our digital lives; we just need to trust that groups and communities are ready to tell their story.

“There is something in us that keeps us where we find ourselves. I think this is the most awful thing of all” (Pyotr Ouspenskii).

References

- AARP. (2020, May). Caregiving in the US. Retrieved 11 May 2022 from <https://www.aarp.org/content/dam/aarp/ppi/2020/05/full-report-caregiving-in-the-united-states.doi.10.26419-2Fppi.00103.001.pdf>.
- Abirami, B., Subashini, T., & Mahavaishnavi, V. (2020). Gender and age prediction from real time facial images using CNN. *Materials Today: Proceedings*, 33, 4708–4712. Retrieved from <https://www.sciencedirect.com/science/article/pii/S2214785320362222> (International Conference on Nanotechnology: Ideas, Innovation and Industries) doi: 10.1016/j.matpr.2020.08.350
- Agarwal, S., & Mishra, S. (2021). *Responsible AI : implementing ethical and unbiased algorithms*. Springer.
- Aitken, M., de Jorre, J. S., Pagliari, C., Jepson, R., & Cunningham-Burley, S. (2016). Public responses to the sharing and linkage of health data for research purposes: a systematic review and thematic synthesis of qualitative studies. *BMC Medical Ethics*, 17(73). doi: 10.1186/s12910-016-0153-x
- Amoore, L. (2018). Cloud geographies: Computing, data, sovereignty. *Progress in Human Geography*, 42(1), 4–24.
- Anderson, T., & Quach, K. (2021, July 6). GitHub Copilot auto-coder snags emerge, from seemingly spilled secrets to bad code, but some love it. *The Register*. Retrieved 5 May 2022 from https://www.theregister.com/2021/07/06/github_copilot_autocoder_caught_spilling/.
- Baran, P. (1964, August). On distributed communications: Introduction to distributed communications networks. *RAND Corporation*. Retrieved from https://www.rand.org/content/dam/rand/pubs/research_memoranda/2006/RM3420.pdf
- Bartlett, M. (2021). Beyond Privacy: Protecting Data Interests in the Age of Artificial Intelligence. *Law, Technology and Humans*, 3(1), 96–108.
- Baskerville, R., Baiyere, A., Gregor, S., Hevner, A., Rossi, M., et al. (2018). Design Science Research contributions: Finding a balance between artifact and theory. *Journal of the Association for Information Systems*, 19(5), 358–376.
- Bayless, T. (1996). The preservation of good medicine is dependant on information. *Journal of*

References

- the Florida Medical Association*, 83(9), 639–642.
- Belfrage, S., Helgesson, G., & Lynøe, N. (2022). Trust and digital privacy in healthcare: a cross-sectional descriptive study of trust and attitudes towards uses of electronic health data among the general public in Sweden. *BMC Medical Ethics*, 23(19), 1–8.
- Beltadze, D. (2020). Developing methodology for the register-based census in Estonia. *Statistical Journal of the IAOS*, 36(2020), 159–164.
- Benbasa, I., & Zmud, R. (1999). Empirical research in information systems: The practice of relevance. *MIS Quarterly*, 23(1), 3–16. doi: 10.2307/249403
- Biddle, S., & Poulson, J. (2022, 22 April). American phone-tracking firm de-moed surveillance powers by spying on CIA and NSA. *The Intercept*. Retrieved 24 April from <https://theintercept.com/2022/04/22/anomaly-six-phone-tracking-signal-surveillance-cia-nsa/>.
- Bierema, A., Hoskinson, A.-M., Moscarella, R., Lyford, A., Haudek, K., Merrill, J., & Urban-Lurain, M. (2021). Quantifying cognitive bias in educational researchers. *International Journal of Research and Method in Education*, 44(4), 395–413.
- Bishop, D. (2016). Indigenous peoples and the official statistics system in Aotearoa/New Zealand. In T. Kukutai & J. Taylor (Eds.), *Indigenous Data Sovereignty: Toward an agenda* (pp. 291–306). Canberra, Australia: ANU Press.
- Biswas, S. (2017, August 24). How significant is India’s landmark privacy judgement? *BBC*. Retrieved 11 May 2022 from <http://www.bbc.com/news/world-asia-india-41037992>.
- Biswas, S. (2018, March 27). Aadhaar: Is India’s biometric ID scheme hurting the poor? *BBC*. Retrieved 11 May 2022 from <http://www.bbc.com/news/world-asia-india-43207964>.
- Bollinger, J., Zuk, P., Majumder, M., Versalovic, E., Villanueva, A., Hsu, R., . . . Cook-Deegan, R. (2019). What is a Medical Information Commons? *The Journal of Law, Medicine and Ethics*, 47, 41–50.
- Booch, G. (2015). Foreword. In G. Suryanarayana, G. Samarthiyam, & T. Sharma (Eds.), *Refactoring for Software Design Smells: Managing Technical Debt*. Waltham, MA, USA: Elsevier.
- Boom, D. V. (2022, May 5). A Typo Sent \$36 Million of Crypto Into the Ether. *CNET*. Retrieved 7 May 2022 from <https://www.cnet.com/personal-finance/crypto/a-typo-sent-36-million-of-crypto-into-the-ether/>.
- Bordoloi, P., & Islam, N. (2012). Knowledge Management practices and healthcare delivery: A contingency framework. *The Electronic Journal of Knowledge Management*, 10(2), 110–120.

References

- Bowman, C. (2017, January 6). Data Localization Laws: an Emerging Global Trend. *Jurist*. Retrieved 11 May 2022 from <http://www.jurist.org/hotline/2017/01/Courtney-Bowman-data-localization.php>.
- Bradshaw, S., Millard, C., & Walden, I. (2011). Contracts for Clouds: Comparison and analysis of the terms and conditions of cloud computing services. *International Journal of Law and Information Technology*, 19(3), 187–223.
- Branscombe, M. (2016, 12 July). Stop saying the cloud is just someone else's computer - because it's not. *ZDNet*. Retrieved 16 April from <https://www.zdnet.com/article/stop-saying-the-cloud-is-just-someone-elses-computer-because-its-not/>.
- Braun, V. (2019, April). Answers to frequently asked questions about thematic analysis. Retrieved 11 May 2022 from <https://cdn.auckland.ac.nz/assets/psych/about/our-research/documents/Answers%20to%20frequently%20asked%20questions%20about%20thematic%20analysis%20April%202019.pdf>.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.
- Burns, D., Hambleton, R., & Hoggett, P. (1994). Neighbourhood Decentralisation and the New Public Management. In *The Politics of Decentralisation: Revitalising Local Democracy* (pp. 81–110). London: Macmillan Education UK.
- Buterin, V. (2016). Hard Fork Completed. *Ethereum Foundation*. Retrieved 5 May 2022 from <https://blog.ethereum.org/2016/07/20/hard-fork-completed/>.
- Castillo-Montoya, M. (2016). Preparing for Interview Research: The Interview Protocol Refinement Framework. *The Qualitative Report*, 21(5), 811–831.
- Chacon, S., & Straub, B. (2020). *Pro Git* (2nd ed.). Apress.
- Chandel, S., Jingji, Z., Yunnan, Y., Jingyao, S., Zhipeng, Z., et al. (2019, October 17-19). The Golden Shield Project of China: A Decade Later—An in-Depth Study of the Great Firewall. In *Proceedings of the 2019 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC)* (pp. 111–119). Giulini, China.
- Chatterjee, S., & Hevner, A. (2010). *Design research in information systems : theory and practice*. Boston, MA, USA: Springer.
- Choi, M., Budak, C., Romero, D. M., & Jurgens, D. (2021, June 7-10). More than Meets the Tie: Examining the Role of Interpersonal Relationships in Social Networks. In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media (ICWSM2021)* (pp. 105–116).
- Choudhary, M. (2018, April 23). Viewpoint: The pitfalls of India's biometric ID scheme. *BBC*. Retrieved 11 May 2022 from <http://www.bbc.com/news/world-asia-india-43619944>.

References

- Citron, D. (2021). A new compact for sexual privacy. *William & Mary Law Review*, 62(6), 1763–1840.
- Clemence, M. (2021, 12 October). Doctors become the world's most trusted profession. *Ipsos*. Retrieved 11 May 2022 from <https://www.ipsos.com/en/global-trustworthiness-index-2021>.
- Codd, E. (1970, June). A relational model of data for large shared data banks. *Communications of the ACM*, 13(6), 377–387.
- Coffey, A., & Atkinson, P. (1996). *Making sense of Qualitative data: Complementary research strategies*. Thousand Oaks, CA, USA: Sage.
- Coiera, E. (2009). Building a National Health IT System from the Middle Out. *Journal of the American Medical Informatics Association*, 16(3), 271–273.
- Collaboration on International ICT Policy for East and Southern Africa (CIPESA), et al. (2019). The Highs and Lows of Uganda's Data Protection and Privacy Act 2019. Retrieved 26 April 2022 from https://www.cipesa.org/?wpfb_dl=303.
- Conger, K. (2016, November 16). China's new cybersecurity law is bad news for business. *TechCrunch*. Retrieved 11 May 2022 from <https://techcrunch.com/2016/11/06/chinas-new-cybersecurity-law-is-bad-news-for-business/>.
- Cowles, E., & Nelson, E. (2015). *An Introduction to Survey Research*. Business Expert Press.
- Croman, K., Decker, C., Eyal, I., Gencer, A. E., Juels, A., Kosba, A., ... Wattenhofer, R. (2016). On Scaling Decentralized Blockchains. In J. Clark, S. Meiklejohn, P. Y. Ryan, D. Wallach, M. Brenner, & K. Rohloff (Eds.), *Financial Cryptography and Data Security* (pp. 106–125). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Cyphers, B., & Gebhart, G. (2019, December 2). Behind the One-Way Mirror: A Deep Dive Into the Technology of Corporate Surveillance. *Electronic Frontier Foundation*. Retrieved 11 May 2022 from <https://www.eff.org/wp/behind-the-one-way-mirror>.
- Dailey, N. (2022, 21 March). Crypto isn't decentralized. It's actually run by a handful of big wigs exploiting low-paid workers, says long-time internet academic. *Business Insider*. Retrieved 7 April 2022 from <https://markets.businessinsider.com/news/currencies/crypto-isnt-decentralized-nft-bored-ape-yacht-club-buys-cryptopunks-2022-3>.
- Dass, R. (2011). Unique Identification for Indians: A Divine Dream or a Miscalculated Heroism? *Vikalpa*, 36(1), 1–14.
- Data Futures Partnership. (2017a). Our data, our way - What New Zealand people expect from guidelines for data use and sharing. Retrieved 21 April 2019 from <http://www.datafutures.co.nz/assets/Uploads/Our-Data-Our-Way-Final-Report.pdf>.
- Data Futures Partnership. (2017b). A Path to Social Licence. Retrieved 21 April

References

- 2019 from <https://trusteddata.co.nz/wp-content/uploads/2017/08/Summary-Guidelines.pdf>. (Retrieved from <https://trusteddata.co.nz/wp-content/uploads/2017/08/Summary-Guidelines.pdf>)
- Dattani, K. (2020). "Goventrepreneurism" for good governance: The case of Aadhaar and the India Stack. *Area*, 52(2), 411–419.
- Davis, M. (2016). Data and the United Nations Declaration on the Rights of Indigenous Peoples. In T. Kukutai & J. Taylor (Eds.), *Indigenous Data Sovereignty: Toward an agenda* (pp. 25–38). Canberra, Australia: ANU Press.
- Dawson, P. (2020). *What are the barriers to equitable maternal health in Aotearoa New Zealand?* (Doctoral dissertation, University of Otago). Retrieved 11 May 2022 from <https://ourarchive.otago.ac.nz/bitstream/handle/10523/10786/P%20DAWSON%20PhD%204401613.pdf>.
- de Aguiar, E., Faical, B., Krishnamachari, B., & Ueyama, J. (2020). A Survey of Blockchain-Based Strategies for Healthcare. *ACM Computing Surveys*, 53(2), 1–27.
- Debreceeny, R., Putterill, M., Tung, L., Gilbert, A., et al. (2002). New tools for the determination of e-commerce inhibitors. *Decision Support Systems*, 34(2), 177–195.
- Deepalakshmi, K. (2017, March 24). The long list of Aadhaar-linked schemes. *The Hindu*. Retrieved 11 May 2022 from <http://www.thehindu.com/news/national/the-long-list-of-aadhaar-linked-schemes/article17641068.ece>.
- de Montjoye, Y.-A., Hidalgo, C., Verleysen, M., Blondel, V., et al. (2013). Unique in the Crowd: The privacy bounds of human mobility. *Scientific Reports*, 3(1376). doi: 10.1038/srep01376
- Desai, S., & Jasuja, N. (2016, October 28). India Stack: The Bedrock of a Digital India. *Medium*. Retrieved 28 April 2022 from <https://medium.com/wharton-fintech/the-bedrock-of-a-digital-india-3e96240b3718>.
- de Tocqueville, A. (1898). *Democracy in America*.
- digital.govt.nz, et al. (2020, July 1). International Open Data Charter. Retrieved 7 May 2022 from <https://www.digital.govt.nz/digital-government/international-partnerships/international-open-data-charter/>.
- Ding, Y., & Sato, H. (2020). Derepo: A Distributed Privacy-Preserving Data Repository with Decentralized Access Control for Smart Health. In *2020 7th IEEE International Conference on Cyber Security and Cloud Computing (CSCloud)/2020 6th IEEE International Conference on Edge Computing and Scalable Cloud (EdgeCom)* (pp. 29–35). doi: 10.1109/CSCloud-EdgeCom49738.2020.00015
- Doberstein, C., Charbonneau, E., Morin, G., Despatie, S., et al. (2022). Measuring the Acceptability of Facial Recognition-Enabled Work Surveillance Cameras in the Public and

References

- Private Sector. *Public Performance & Management Review*, 45(1), 198–227.
- Dreze, J., Khalid, N., Khera, R., Somanchi, A., et al. (2017). Aadhaar and food security in Jharkhand pain without gain? *Economic and Political Weekly*, 52, 50–60.
- Dubovitskaya, A., Xu, Z., Ryu, S., Schumacher, M., Wang, F., et al. (2017). Secure and Trustable Electronic Medical Records Sharing using Blockchain. *AMIA ... Annual Symposium proceedings. AMIA Symposium, 2017*, 650—659. Retrieved 11 May 2022 from <https://europepmc.org/articles/PMC5977675>.
- E-Estonia. (2017, June). Estonia to open the world's first data embassy in Luxembourg. *E-Estonia*. Retrieved 11 May 2022 from <https://e-estonia.com/estonia-to-open-the-worlds-first-data-embassy-in-luxembourg/>.
- Edwards, N., Kornacki, M., & Silversin, J. (2002). Unhappy doctors: What are the causes and what can be done? *British Medical Journal*, 324(7341), 177–195.
- Eichstaedt, J., & Weidman, A. (2020). Tracking Fluctuations in Psychological States using Social Media Language: A Case Study of Weekly Emotion. *European Journal of Personality*, 34(5), 845–858.
- Ekblaw, A., Azaria, A., Halamka, J., & Lippmann, A. (2016, August 22-24). A Case Study for Blockchain in Healthcare: “MedRec” prototype for electronic health records and medical research data. In *Proceedings of the 2016 IEEE Open & Big Data Conference. 2016 Presented at: IEEE BigData'16*. Washington DC, USA.
- Estonia Ministry of Economic Affairs and Communications. (2016). Implementation of the Virtual Data Embassy Solution. Retrieved 11 May 2022 from <https://www.microsoft.com/en-us/cybersecurity/content-hub/implementation-of-the-virtual-data-embassy-solution>.
- European Commission. (2013, October 15). What does the Commission mean by secure Cloud computing services in Europe? Retrieved 11 May 2022 from http://europa.eu/rapid/press-release_MEMO-13-898_en.htm.
- European Commission, Directorate-General of Communications Networks, Content & Technology. (2014, April 8). Study on eGovernment and the Reduction of Administrative Burden. *European Commission*. Retrieved 14 April 2022 from <https://digital-strategy.ec.europa.eu/en/library/final-report-study-egovernment-and-reduction-administrative-burden-smart-20120061>.
- Eyal, I., & Sirer, E. G. (2014). Majority Is Not Enough: Bitcoin Mining Is Vulnerable. In N. Christin & R. Safavi-Naini (Eds.), *Financial Cryptography and Data Security* (pp. 436–454). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Fink, A. (2003). *The Survey kit: How to sample in surveys*. Thousand Oaks, CA, USA: Sage.
- Foster, J., & Parker, I. (1995). *Carrying out investigations in psychology: methods and statistics*.

References

- Leicester, UK: BPS Books.
- Frahat, R., Monowar, M., & Buhari, S. (2019, 05). Secure and Scalable Trust Management Model for IoT P2P Network. In (pp. 1–6). doi: 10.1109/CAIS.2019.8769467
- Gallagher, S. (2019, October). 50 years ago today, the Internet was born. Sort of. *Ars Technica*. Retrieved 11 May 2022 from <https://arstechnica.com/information-technology/2019/10/50-years-ago-today-the-internet-was-born-sort-of>.
- Garett, R., & Young, S. D. (2022). Ethical Views on Sharing Digital Data for Public Health Surveillance: Analysis of Survey Data Among Patients. *Frontiers in Big Data*, 5. doi: 10.3389/fdata.2022.871236
- Garrison, N., Sathe, N., Antommaria, A., Holm, I., Sanderson, S., Smith, M., ... others (2016). A systematic literature review of individuals' perspectives on broad consent and data sharing in the United States. *Official Journal of the American College of Medical Genetics and Genomics*, 18(7), 663–671.
- Garvie, C. (2020, 24 June). The Untold Number of People Implicated in Crimes They Didn't Commit Because of Face Recognition. *American Civil Liberties Union (ACLU)*. Retrieved 14 April from <https://www.aclu.org/news/privacy-technology/the-untold-number-of-people-implicated-in-crimes-they-didnt-commit-because-of-face-recognition>.
- Gauld, R. (2004). One step forward, one step back? Restructuring, evolving policy, and information management and technology in the New Zealand health sector. *Government Information Quarterly*, 21(2), 125–142.
- Gazem, N., Rahman, A., Saeed, F., & Iahad, N. (2018). Design Science Research Roadmap Model for Information Systems Projects: A Case Study. *International Journal of Information Technology Project Management*, 9(3). doi: 10.4018/IJITPM.2018070101
- Geekiyange Don, J. S., & Motalebi, F. (2021). Decentralization using Blockchain Health Records Management. In *2021 International Conference on Green Energy, Computing and Sustainable Technology (GECOST)* (pp. 1–5). doi: 10.1109/GECOST52368.2021.9538734
- Gershgorin, D. (2021, July 7). GitHub's automatic coding tool rests on untested legal ground. *The Verge*. Retrieved 5 May 2022 from <https://www.theverge.com/2021/7/7/22561180/github-copilot-legal-copyright-fair-use-public-code>.
- Gong, V., Daamen, W., Bozzon, A., & Hoogendoorn, S. (2021). Counting people in the crowd using social media images for crowd management in city events. *Transportation*, 48, 3085–3119.
- Gottsegen, W. (2021, 19 October). The NFT Market Is Already Centralized. *CoinDesk*. Retrieved 22 April from <https://www.coindesk.com/tech/2021/10/18/the>

References

- [-nft-market-is-already-centralized/](#).
- Graber, J. (2021, January). Ecosystem review. *Twitter Bluesky*. Retrieved 11 May 2022 from https://matrix.org/_matrix/media/r0/download/twitter.modular.im/981b258141aa0b197804127cd2f7d298757bad20.
- Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360–1380.
- Greenhalgh, T., Stramer, K., Bratan, T., Byrne, E., Russell, J., Hinder, S., & Potts, H. (2010). *The devil's in the detail: Final report of the independent evaluation of the Summary Care Record and HealthSpace programmes*. University College London.
- Gregor, S., & Hevner, A. (2013). Positioning and presenting Design Science Research for maximum impact. *MIS Quarterly*, 37(2), 337–355.
- Gregor, S., & Jones, D. (2007). The anatomy of a design theory. *Journal of the Association for Information Systems*, 8(5), 312–335. doi: 10.17705/1jais.00129
- Groves, R., Fowler, F., Jr, Couper, M., Lepkowski, J., Singer, E., & Tourangeau, R. (2009). *Survey Methodology* (2nd ed.). Hoboken, NJ, USA: Wiley.
- Gunter, J. (2021, 9 December). China committed genocide against Uyghurs, independent tribunal rules. *BBC News*. Retrieved 24 April from <https://www.bbc.com/news/world-asia-china-59595952>.
- Hartz, S. M., Johnson, E. O., Saccone, N. L., Hatsukami, D., Breslau, N., & Bierut, L. J. (2011). *Practice of Epidemiology Inclusion of African Americans in Genetic Studies: What Is the Barrier?* (Vol. 174) (No. 3).
- Harvey, D. (2016). Privacy and new technologies. In S. Penk & R. Tobin (Eds.), *Privacy law in New Zealand* (2nd ed., pp. 385–428). Wellington, New Zealand: Thomson Reuters New Zealand.
- Hate, K., Meherally, S., More, N., Jayaraman, A., Bull, S., Parker, M., & Osrin, D. (2015). Sweat, Skepticism, and Uncharted Territory: A Qualitative Study of Opinions on Data Sharing Among Public Health Researchers and Research Participants in Mumbai, India. *Journal of Empirical Research on Human Research Ethics*, 10(3), 239–250. doi: 10.1177/1556264615592383
- Health and Disability System Review. (2020). Health and Disability System Review – Final Report – Pūrongo Whakamutunga. Retrieved 11 May 2022 from <https://systemreview.health.govt.nz/assets/Uploads/hdsr/health-disability-system-review-final-report.pdf>.
- Health Quality & Safety Commission. (2019, May). A window on the quality of Aotearoa New Zealand's Health Care 2019. Retrieved 11 May 2022 from <https://www.hqsc.govt.nz/assets/Our-data/Publications-resources/>

References

- [Window_2019_web_final-v2.pdf](#).
- Heller, N. (2017, December 18). Estonia, the Digital Republic. *The New Yorker*. Retrieved 11 May 2022 from <https://www.newyorker.com/magazine/2017/12/18/estonia-the-digital-republic>.
- Hemel, A., & Coughlan, S. (2017). Making Sense Of Git In A Legal Context. *International Free and Open Source Software Law Review*, 9(1), 19–34.
- Henry, J., Pylypchuk, Y., Searcy, Y., Patel, V., et al. (2016, May). Adoption of Electronic Health Record Systems among U.S. Non-Federal Acute Care Hospitals: 2008-2015. *The Office of the National Coordinator for Health Information Technology*. Retrieved 9 April 2022 from https://www.healthit.gov/sites/default/files/briefs/2015_hospital_adoption_db_v17.pdf.
- Herlihy, P. (2013, October 31). ‘Government as a data model’ : what I learned in Estonia. *Government Digital Service*. Retrieved 11 April 2022 from <https://gds.blog.gov.uk/2013/10/31/government-as-a-data-model-what-i-learned-in-estonia/>.
- Herskind, L., Katsikouli, P., & Dragoni, N. (2020). Privacy and Cryptocurrencies—A Systematic Literature Review. *IEEE Access*, 8, 54044–54049.
- Hertzum, M. (2020). *Usability Testing : A Practitioner’s Guide to Evaluating the User Experience*. Morgan & Claypool.
- Herzberg, R. (2020). Elinor Ostrom’s Governing the Commons: Institutional Diversity, Self-Governance, and Tragedy Diverted. *The Independent Review*, 24(4), 627–636.
- Hess, C., & Ostrom, E. (2011). Introduction: An overview of the knowledge commons. In C. Hess & E. Ostrom (Eds.), *Understanding knowledge as a commons: From theory to practice* (pp. 3–26). Cambridge, MA, USA: MIT Press.
- Hevner, A., March, S., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, 28(1), 75–105.
- Hobbes, T. (1962). *Leviathan: or, The matter, forme and power of a commonwealth, ecclesiasticall and civil*. New York, NY: Collier Books.
- Hoffman, M., Ibáñez, L.-D., & Simper, E. (2020). Toward a Formal Scholarly Understanding of Blockchain-Mediated Decentralization: A Systematic Review and a Framework. *Frontiers in Blockchain*, 3.
- Holotescu, V., & Vasiu, R. (2020, 03). Challenges and Emerging Solutions for Public Blockchains. *BRAIN. Broad Research in Artificial Intelligence and Neuroscience*, 11, 58–83. doi: 10.18662/brain/11.1/15
- Hoofnagle, C., & King, J. (2008). Research Report: What Californians Understand About Privacy Online. *University of California, Berkeley, School of Law*. Retrieved 11 May 2022 from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1262130.

References

- Human Rights Watch. (2016, November 6). *China: Abusive Cybersecurity Law Set to be Passed*. Retrieved 11 May 2022 from <https://www.hrw.org/news/2016/11/06/china-abusive-cybersecurity-law-set-be-passed>.
- Hunt, A., & Thomas, D. (2000). *The Pragmatic Programmer: From Journeyman to Master*. Reading, Mass.: Addison-Wesley.
- Hurley, D. (2018, January 2). Can an Algorithm Tell When Kids Are in Danger? *The New York Times*. Retrieved 11 May 2022 from <https://www.nytimes.com/2018/01/02/magazine/can-an-algorithm-tell-when-kids-are-in-danger.html>.
- Hurst, S., Arulogun, O., Owolabi, M., Akinyemi, R., Uvere, E., Warth, S., & Ovbiagele, B. (2015). Pretesting qualitative data collection procedures to facilitate methodological adherence and team building in Nigeria. *International Journal of Qualitative Methods*, 14, 53–64.
- Iivari, J., & Venable, J. (2009). *Action research and design science research - Seemingly similar but decisively dissimilar*. Paper presented at the European Conference on Information Systems (2009). Italy. Retrieved from <https://aisel.aisnet.org/ecis2009/255>
- IPFS. (2017, May 4). Uncensorable Wikipedia on IPFS. Retrieved 7 May 2022 from <https://blog.ipfs.io/24-uncensorable-wikipedia/>.
- Irion, K. (2012). Government Cloud Computing and National Data Sovereignty. *Policy and Internet*, 4(3-4), 40–71.
- Islam, T., & Goldwasser, D. (2021, June 7-10). Analysis of Twitter Users' Lifestyle Choices using Joint Embedding Model. In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media (ICWSM2021)* (pp. 242–253).
- Janjua, K., Shah, M., Almogren, A., Khattak, H., Maple, C., & Din, I. (2020). Proactive forensics in IoT: Privacy-aware log-preservation architecture in fog-enabled-cloud using holochain and containerization technologies. *Electronics*, 9(7), 1–39.
- Jansen, R. (2016). Indigenous data sovereignty: a Maori health perspective. In T. Kukutai & J. Taylor (Eds.), *Indigenous Data Sovereignty: Toward an agenda* (pp. 193–211). Canberra, Australia: ANU Press.
- Jao, I., Kombe, F., Mwalukore, S., Bull, S., Parker, M., Kamuya, D., . . . Marsh, V. (2015). Involving Research Stakeholders in Developing Policy on Sharing Public Health Research Data in Kenya: Views on Fair Process for Informed Consent, Access Oversight, and Community Engagement. *Journal of Empirical Research on Human Research Ethics*, 10(3), 264–277. doi: 10.1177/1556264615592385
- Jones, S., Torres, V., & Arminio, J. (2013). *Negotiating the complexities of qualitative research in higher education: Fundamental elements and issues* (2nd ed.). New York: Routledge.
- Joseph, K. (2021, January 29). Police Needs Sh15Billion to Set up Crim-

References

- inal Identification Systems. *Uganda Radio Network*. Retrieved 28 April 2022 from <https://ugandaradionetwork.net/story/police-needs-sh15billion-to-set-up-criminal-identification-systems>.
- Kafeero, S. (2020, November 28). Uganda is using Huawei's facial recognition tech to crack down on dissent after anti-government protests. *Quartz Africa*. Retrieved 28 April 2022 from <https://qz.com/africa/1938976/uganda-uses-chinas-huawei-facial-recognition-to-snare-protesters/>.
- Kanat, O. (2021). The Uyghur Genocide and International Policy Response. *Brown Journal of World Affairs*, 28(1), 1–15.
- Kang, C., & Frenkel, S. (2018, April 4). Facebook Says Cambridge Analytica Harvested Data of Up to 87 Million Users. *The New York Times*. Retrieved 11 May 2022 from <https://www.nytimes.com/2018/04/04/technology/mark-zuckerberg-testify-congress.html>.
- Kass, N., Taylor, H., Ali, J., Hallez, K., Chaisson, L., et al. (2015). A pilot study of simple interventions to improve informed consent in clinical research: Feasibility, approach, and results. *Clinical Trials*, 12(1), 54–66. doi: 10.1177/1740774514560831
- Khawwada, S., Tushev, M., & Mahmoud, A. (2018). Just enough semantics: An information theoretic approach for IR-based software bug localization. *Information and Software Technology*, 93, 45–57.
- Khera, R. (2017, December 16). Impact of Aadhaar on Welfare Programmes. *Economic and Political Weekly*, 52(50).
- Kinsbruner, E., & Bahmutov, G. (2022). *A Frontend Web Developer's Guide to Testing : Explore Leading Web Test Automation Frameworks and Their Future Driven by Low-Code and AI*. Packt.
- Kirchgaessner, S. (2022, 2 February). FBI confirms it obtained NSO's Pegasus spyware. *The Guardian*. Retrieved 20 April from <https://www.theguardian.com/news/2022/feb/02/fbi-confirms-it-obtained-nsos-pegasus-spyware>.
- Konkel, F. (2014, 18 July). The Details About the CIA's Deal With Amazon. *The Atlantic*. Retrieved 24 April from <https://www.theatlantic.com/technology/archive/2014/07/the-details-about-the-cias-deal-with-amazon/374632/>.
- Koppel, R. (2012, July 1). Patient Safety and Health Information Technology: Learning from Our Mistakes. *Agency for Healthcare Research and Quality*. Retrieved 9 April 2022 from <https://psnet.ahrq.gov/perspective/patient-safety-and-health-information-technology-learning-our-mistakes>.
- Kosinski, M., & Wang, Y. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psy-*

References

- chology*, 114(2), 246–257.
- Krämer, N., Sauer, V., & Ellison, N. (2021, April - June). The Strength of Weak Ties Revisited: Further Evidence of the Role of Strong Ties in the Provision of Online Social Support. *Social Media and Society*, 1–19. doi: 10.1177/20563051211024958
- Krill, P. (2017, March 6). WebAssembly is now ready for browsers to use. *InfoWorld*. Retrieved 15 April 2022 from <https://www.infoworld.com/article/3176681/webassembly-is-now-ready-for-browsers-to-use.html>.
- Krosnick, J., & Fabrigar, L. (1997). Designing rating scales for effective measurement in surveys. In P. Lyberg, M. Collins, E. de Leeuw, C. Dippo, N. Schwarz, & D. Trewin (Eds.), *Survey measurement and process quality* (pp. 141–164). New York, NY: Wiley.
- Kukutai, T., & Cormack, D. (2021). “Pushing the space” - Data sovereignty and self-determination in Aotearoa NZ. In M. Walter, T. Kukutai, S. R. Carroll, & D. Rodriguez-Lonebear (Eds.), *Indigenous Data Sovereignty and Policy* (pp. 21–35). Abingdon, Oxon, UK: Routledge.
- Kukutai, T., & Taylor, J. (Eds.). (2016). *Indigenous Data Sovereignty: Toward an agenda*. Canberra, Australia: ANU Press.
- Kushwaha, N., Roguski, P., & Watson, B. W. (2020). Up in the Air: Ensuring Government Data Sovereignty in the Cloud. In *2020 12th International Conference on Cyber Conflict (CyCon)* (Vol. 1300, pp. 43–61). doi: 10.23919/CyCon49761.2020.9131718
- Lam, O. (2020, 31 March). Leaked Xinjiang police report describes circumvention tools as ‘terrorist software’. *Hong Kong Free Press*. Retrieved 24 April from <https://hongkongfp.com/2016/10/29/leaked-xinjiang-police-report-describes-circumvention-tools-terrorist-software/>.
- Lavrakas, P. (2008). *Encyclopedia of survey research methods*. Thousand Oaks, CA, USA: Sage.
- Lee, H., Kung, H., Udayasankaran, J., Kijisanayotin, B., Marcelo, A., Chao, L., & Hsu, C.-Y. (2020). An Architecture and Management Platform for Blockchain-Based Personal Health Record Exchange: Development and Usability Study. *Journal of Medical Internet Research*, 22(6), e16748. doi: 10.2196/16748
- Lee, J.-A. (2018). Hacking into China’s Cybersecurity Law. *Wake Forest Law Review*, 53(1), 57–104.
- Lee, L., Lee, J., Egelman, S., & Wagner, D. (2016, 01). Information Disclosure Concerns in The Age of Wearable Computing.. doi: 10.14722/usec.2016.23006
- Lepore, J. (2013, June 24). The prism: Privacy in an age of publicity. *The New Yorker*. Retrieved 11 May 2022 from <https://www.newyorker.com/magazine/2013/06/24/the-prism>.

References

- Levine, P. (2011). Collective action, civic engagement, and the knowledge commons. In C. Hess & E. Ostrom (Eds.), *Understanding knowledge as a commons: From theory to practice* (pp. 247–276). Cambridge, MA, USA: MIT Press.
- Lewis, C. (1982). Using the "Thinking-aloud" Method in Cognitive Interface Design. *IBM Thomas Watson Research Center*. Retrieved 8 April 2022 from <https://dominoweb.draco.res.ibm.com/reports/RC9265.pdf>.
- Li, W., He, M., & Haiquan, S. (2021). An Overview of Blockchain Technology: Applications, Challenges and Future Trends. In *2021 IEEE 11th International Conference on Electronics Information and Emergency Communication (ICEIEC)* (pp. 31–39). doi: 10.1109/ICEIEC51955.2021.9463842
- Lieb, D. (2017, June 26). Analysis indicates partisan gerrymandering has benefited GOP. *Associated Press*. Retrieved 11 May 2022 from <https://apnews.com/fa6478e10cda4e9cbd75380e705bd380>.
- Lindsay, J. (2014). The Impact of China on Cybersecurity: Fiction and Friction. *International Security*, 39(3), 7–47.
- Lunden, I. (2014, December 6). To Get Off Russia's Blacklist, GitHub Has Blocked Access To Pages That Highlight Suicide. *TechCrunch*. Retrieved 11 May 2022 from <https://techcrunch.com/2014/12/05/to-get-off-russias-blacklist-github-has-blocked-access-to-pages-that-highlight-suicide>.
- Mackieson, P., Shlonsky, A., & Connolly, M. (2019). Increasing rigor and reducing bias in qualitative research: A document analysis of parliamentary debates using applied thematic analysis. *Qualitative Social Work*, 18(6), 965–980.
- Maier, M., & Harr, R. (2020). Dark design patterns: An end-user perspective. *Human Technology*, 16(2), 170–199.
- Malgieri, G., & Comandé, G. (2017). Why a Right to Legibility of Automated Decision-Making Exists in the General Data Protection Regulation. *International Data Privacy Law*, 7(4), 243–265. Retrieved from <https://doi.org/10.1093/idpl/ix019>
- Malkin, N., Bernd, J., Johnson, M., Egelman, S., et al. (2018, April). "What can't data be used for?" *Privacy expectations about smart TVs in the US*. Paper presented at the European Workshop on Usable Security (EuroUSEC) 2018. London, England. Retrieved from <https://dx.doi.org/10.14722/eurosec.2018.23007>
- Malterud, K., Siersma, V. D., & Guassora, A. D. (2016). Sample Size in Qualitative Interview Studies: Guided by Information Power. *Qualitative Health Research*, 26(13), 1753–1760.
- Mandavia, M. (2019, December 12). Personal Data Protection Bill can turn India into 'Orwellian State': Justice BN Srikrishna. *The Economic Times*. Retrieved 14 April

References

- 2022 from <https://economictimes.indiatimes.com/news/economy/policy/personal-data-protection-bill-can-turn-india-into-orwellian-state-justice-bn-srikrishna/articleshow/72483355.cms>.
- Mandl, K., & Kohane, I. (2012). Escaping the EHR trap - the future of health IT. *New England Journal of Medicine*, 366(24), 2240–2.
- March, S., & Smith, G. (1995). Design and Natural Science Research on Information Technology. *Decision Support Systems*, 15(4), 251–266.
- Martindale, S., & Taylor, R. (2014). Alternative and complementary research approaches. In R. Taylor (Ed.), *The Essentials of Nursing and Healthcare Research*. Thousand Oaks, CA, USA: Sage.
- Mason, M. (2010). Sample size and saturation in PhD studies using qualitative interviews. *Forum: Qualitative Social Research*, 11(8).
- Maxwell, J. (2012). *Qualitative Research Design: An Interactive Approach* (3rd ed.). Thousand Oaks, CA, USA: Sage.
- McCluskey, M. (2015, March 20). Estonia redefines national security in a digital age. *Al Jazeera*. Retrieved 11 May 2022 from <https://www.aljazeera.com/indepth/features/2015/03/estonia-redefines-national-security-digital-age-150318065430514.html>.
- McMillan, R. (2012, November 5). Linus Torvalds Invented Git, But He Pulls No Patches With GitHub. *Wired*. Retrieved 11 May 2022 from <https://www.wired.com/2012/05/torvalds-github>.
- McMullan, T. (2015, October 18). Guardian readers on privacy: "we trust government over corporations". *The Guardian*. Retrieved 11 May 2022 from <https://www.theguardian.com/technology/2015/oct/18/guardian-readers-on-privacy-we-trust-government-over-corporations>.
- MedTech Global. (2021, April 8). The ALEX Platform explained. Retrieved 9 April 2022 from <https://medtechglobal.com/alex-platform-a-game-changer/>.
- Mezuk, B., Eaton, W., & Zandi, P. (2008). Participant characteristics that influence consent for genetic research in a population-based survey: the Baltimore epidemiologic catchment area follow-up. *Community Genetics*, 11(3), 171–178. Retrieved from <http://hdl.handle.net/10822/549376>
- Microsoft. (2017, December 14). Diplomatic immunity for data: Estonia creates a virtual embassy. Retrieved 11 May 2022 from <https://blogs.microsoft.com/eupolicy/2017/12/14/diplomatic-immunity-data-estonia-creates-virtual-embassy/>.
- Microsoft. (2018, June). Microsoft to acquire GitHub for \$7.5 billion. Retrieved 11 May 2022 from <https://news.microsoft.com/2018/06/04/microsoft-to-acquire-github>

References

- for-7-5-billion.
- Miller, C. (2020, September 27). How Taiwan's 'civic hackers' helped find a new way to run the country. *The Guardian*. Retrieved 7 May 2022 from <https://www.theguardian.com/world/2020/sep/27/taiwan-civic-hackers-polis-consensus-social-media-platform>.
- Miller, M., Yee, K.-P., & Shapiro, J. (2003). Capability Myths Demolished. *Technical Report SRL2003-02, Systems Research Laboratory, Johns Hopkins University*. Retrieved 8 April 2022 from <https://srl.cs.jhu.edu/pubs/SRL2003-02.pdf>.
- Milmo, D. (2022, 4 March). Russia blocks access to Facebook and Twitter. *The Guardian*. Retrieved 18 April from <https://www.theguardian.com/world/2022/mar/04/russia-completely-blocks-access-to-facebook-and-twitter>.
- Minevich, M. (2021, 18 January). How The EU Is Leading The Way In AI Powered Social Innovation. *Forbes*. Retrieved 16 April from <https://www.forbes.com/sites/markminevich/2021/01/18/how-the-eu-is-leading-the-way-in-ai-powered-social-innovation>.
- Ministry of Health. (2016, April). New Zealand Health Strategy: Future Direction. Retrieved 11 May 2022 from <https://www.health.govt.nz/system/files/documents/publications/new-zealand-health-strategy-futuredirection-2016-apr16.pdf>.
- Ministry of Health. (2019, March 14). Population of Waikato DHB. Retrieved 11 May 2022 from <https://www.health.govt.nz/new-zealand-health-system/my-dhb/waikato-dhb/population-waikato-dhb>.
- Ministry of Health. (2021a, 16 December). Hira (National health information platform). Retrieved 11 May 2022 from <https://www.health.govt.nz/our-work/digital-health/other-digital-health-initiatives/hira-national-health-information-platform>.
- Ministry of Health. (2021b, 26 February). Hira Programme Business Case. Retrieved 11 May 2022 from https://www.health.govt.nz/system/files/documents/pages/hira_programme_business_case.pdf.
- Ministry of Health. (2022, 17 March). New Zealand ePrescription Service. Retrieved on 3 April 2022 from <https://www.health.govt.nz/our-work/digital-health/other-digital-health-initiatives/emedicines/new-zealand-eprescription-service>.
- Moana Research. (2021, February). Pacific Data Sovereignty Network: Consultation document. Retrieved 11 May 2022 from <https://moanaconnect.co.nz/wp-content/uploads/2021/03/PDS-consultation-document.pdf>.

References

- Monitor Uganda. (2020, September 16). Museveni launches Shs43b data centre in Jinja. Retrieved 28 April 2022 from <https://www.monitor.co.ug/News/National/Museveni-launches-Shs43b-data-centre-Jinja/688334-5192392-p2uf71z/index.html>.
- Moon, L. (2017). Factors influencing health data sharing preferences of consumers: A critical review. *Health Policy and Technology*, 6, 169–187. Retrieved from <http://dx.doi.org/10.1016/j.hlpt.2017.01.001>
- Moore, C. (2019). What works?: Social investment, big data and social services in Aotearoa/New Zealand. *New Zealand Sociology*, 34(2), 123–147.
- Morphy, F. (2016). Indigenising demographic categories: a prolegomenon to indigenous data sovereignty. In T. Kukutai & J. Taylor (Eds.), *Indigenous Data Sovereignty: Toward an agenda* (pp. 99–115). Canberra, Australia: ANU Press.
- Mujuzi, J. (2012). The Right to Privacy of People in or Presumed to Be in Same-Sex Relationships in Uganda. *African Journal of International and Comparative Law*, 20(1), 110–118.
- Murad, A., Myers, M., Thompson, S., Fisher, R., Antommaria, A., et al. (2017). A qualitative study of adolescents' understanding of biobanks and their attitudes toward participation, re-contact, and data sharing. *American Journal of Medical Genetics Part A*, 173(4), 930–937. Retrieved from <https://doi.org/10.1002/ajmg.a.38114>
- Mwaura, W. (2021). *End-To-End Web Testing with Cypress : Explore Techniques for Automated Frontend Web Testing with Cypress and JavaScript*. Packt.
- Naeini, P., Bhagavatula, S., Habib, H., Degeling, M., Bauer, L., Cranor, L., ... others (2017, July 12-14). Privacy Expectations and Preferences in an IoT World. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)* (pp. 399–412). Santa Clara, CA, USA. Retrieved from <https://www.usenix.org/system/files/conference/soups2017/soups2017-naeini.pdf>
- Nair, V. (2021). Becoming data: biometric IDs and the individual in 'Digital India'. *Journal of the Royal Anthropological Institute*, 27, 26–42.
- Nakamoto, S. (2008, 31 October). Bitcoin: A Peer-to-Peer Electronic Cash System. Retrieved 22 April from <https://bitcoin.org/bitcoin.pdf>.
- Nayak, K., Kumar, S., Miller, A., & Shi, E. (2016). Stubborn Mining: Generalizing Selfish Mining and Combining with an Eclipse Attack. In *2016 IEEE European Symposium on Security and Privacy (EuroS P)* (pp. 305–320). doi: 10.1109/EuroSP.2016.32
- Neilson, M. (2021, 11 November). Covid 19 Delta: Ministers at odds with ministry Whānau Ora Māori vax stoush back in Court. *New Zealand Herald*. Retrieved 5 April 2022 from <https://www.nzherald.co.nz/nz/covid-19-delta-ministers-at-odds-with-ministry-whanau-ora-maori-vax-stoush-back-in-court/>

References

- 20KQ7COD7Z52M2DNL52KMISFZM/.
- NIST Big Data Public Working Group Definitions and Taxonomies Subgroup. (2019, October). NIST Big Data Interoperability Framework: Volume 4, Security and Privacy. Retrieved 11 May 2022 from <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-4r2.pdf>.
- Ochiai, K., Fukazawa, Y., Yamada, W., Manabe, H., & Matsuo, Y. (2021, June 7-10). Gravity of Location-Based Service: Analyzing the Effects for Mobility Pattern and Location Prediction. In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media (ICWSM2021)* (pp. 476–487).
- Office of the Registrar General & Census Commissioner, India. (2012). FAQ for NPR. Retrieved 21 April 2019 from <http://www.censusindia.gov.in/2011-Common/FAQs.html>.
- Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57(6), 1701–1777. Retrieved from <https://www.uclalawreview.org/broken-promises-of-privacy-responding-to-the-surprising-failure-of-anonymization-2>
- Oremus, W., Alcantara, C., Merrill, J., Galocha, A., et al. (2021, October 26). How Facebook shapes your feed. *The Washington Post*. Retrieved 5 May 2022 from <https://www.washingtonpost.com/technology/interactive/2021/how-facebook-algorithm-works/>.
- Ormond, K., Cirino, L., Helenowski, I., Chisholm, R., Wolf, W., et al. (2009). Assessing the understanding of biobank participants. *American Journal of Medical Genetics Part A*, 149A, 188–198. Retrieved from <https://doi.org/10.1002/ajmg.a.32635>
- Ostrom, E. (2015). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press. doi: 10.1017/CBO9781316423936
- Osula, A.-M. (2015). Transborder access and territorial sovereignty. *Computer Law and Security Review*, 31, 719–735.
- Pan, S., Stavrinou, T., Zhang, Y., Sikaria, A., Zakharov, P., Sharma, A., ... others (2021, 23 February). *Facebook's Tectonic Filesystem: Efficiency from Exascale*. Proceedings of the 19th USENIX Conference on File and Storage Technologies. Retrieved from <https://www.usenix.org/system/files/fast21-pan.pdf>
- Pandey, G. (2017, August 24). Indian Supreme Court in landmark ruling on privacy. *BBC*. Retrieved 11 May 2022 from <http://www.bbc.com/news/world-asia-india-41033954>.
- Parkinson, J., Bariyo, N., & Chin, J. (2019, August 15). Huawei Technicians Helped African Governments Spy on Political Opponents. *The Wall Street Journal*. Retrieved 28 April 2022 from <https://www.wsj.com/articles/huawei-technicians-helped>

References

- [-african-governments-spy-on-political-opponents-11565793017](#).
- Parrillo, F., & Tschudin, C. (2021). Solid over the Interplanetary File System. In *Proceedings of the 2021 IFIP Networking Conference (IFIP Networking)* (pp. 1–6). doi: 10.23919/IFIPNetworking52078.2021.9472772
- Parsovs, A. (2020, May). Solving the Estonian ID Card Crisis: the Legal Issues. In *Proceedings of the 17th ISCRAM Conference* (pp. 459–471). Blacksburg, VA, USA.
- Patel, V., Barker, W., & Siminerio, E. (2015). Trends in Consumer Access and Use of Electronic Health Information. *Office of the National Coordinator for Health Information Technology*. Retrieved 11 May 2022 from <https://dashboard.healthit.gov/evaluations/data-briefs/trends-consumer-access-use-electronic-health-information.php>.
- Patton, M. (2015). *Qualitative Research and Evaluation Methods* (4th ed.). Thousand Oaks, CA, USA: Sage.
- Pearce, H., Ahmad, B., Tan, B., Dolan-Gavitt, B., Karri, R., et al. (2021). *Asleep at the Keyboard? Assessing the Security of GitHub Copilot’s Code Contributions*. Retrieved 11 May 2022 from <https://arxiv.org/abs/2108.09293>. arXiv.
- Peppers, K., Tuunanen, T., Gengler, C. E., Rossi, M., Hui, W., Virtanen, V., & Bragge, J. (2006). The Design Science Research Process: A Model for Producing and Presenting Information Systems Research. In *In: 1st International Conference on Design Science in Information Systems and Technology (DESRIST)* (pp. 83–106). Claremont, CA, USA.
- Penk, S. (2016a). The Privacy Act 1993. In S. Penk & R. Tobin (Eds.), *Privacy law in New Zealand* (2nd ed., pp. 54–87). Wellington, New Zealand: Thomson Reuters New Zealand.
- Penk, S. (2016b). Thinking about privacy. In S. Penk & R. Tobin (Eds.), *Privacy law in New Zealand* (2nd ed., pp. 1–28). Wellington, New Zealand: Thomson Reuters New Zealand.
- Perez, B., Musolesi, M., & Stringhini, G. (2018). You are your Metadata: Identification and Obfuscation of Social Media Users using Metadata Information. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM 2018)* (pp. 123–132). Stanford, CA, USA. Retrieved from <http://www0.cs.ucl.ac.uk/staff/G.Stringhini/papers/identification-ICWSM2018.pdf>
- Perrin, M. (2017). *Distributed Systems : Concurrency and Consistency*. London, England: ISTE Press.
- Peters, T., & Waterman, R. (1982). *In search of excellence: lessons from America’s best-run companies* (1st ed.). New York: Harper & Row.
- Peterson, K. J., Deeduvanu, R., Kanjamala, P., & Mayo, K. (2016). A Blockchain-Based Approach to Health Information Exchange Networks..
- Petrova, M., Barclay, M., Barclay, S., Barclay, S., et al. (2017). Between “the best way to deliver

References

- patient care” and “chaos and low clinical value”: General Practitioners’ and Practice Managers’ views on data sharing. *International Journal of Medical Informatics*, 104, 74–83. Retrieved from <http://dx.doi.org/10.1016/j.ijmedinf.2017.05.009>
- Pool, I. (2016). Colonialism’s and postcolonialism’s fellow traveller: the collection, use and misuse of data on indigenous people. In T. Kukutai & J. Taylor (Eds.), *Indigenous Data Sovereignty: Toward an agenda* (pp. 57–76). Canberra, Australia: ANU Press.
- Poor, A. (2016). *What are the obstacles to Health Information Systems interoperability in the Auckland region primary health care sector?* (Master’s thesis, Auckland University of Technology). Retrieved from <https://openrepository.aut.ac.nz/bitstream/handle/10292/10190/PoorA.pdf>.
- Porter, J. (2019). GitHub restricts developer accounts based in Iran, Crimea, and other countries under US sanctions. *The Verge*. Retrieved 11 May 2022 from <https://www.theverge.com/2019/7/29/8934694/github-us-trade-sanctions-developers-restricted-crimea-cuba-iran-north-korea-syria>.
- Post, R. (2001). Three concepts of privacy. *Georgetown Law Journal*, 89(6), 2087–2098.
- Privacy Commission. (2017). Inquiry into the Ministry of Social Development’s collection of individual client-level data from NGOs. Retrieved 11 May 2022 from <https://privacy.org.nz/assets/New-order/Resources-/Publications/Commissioner-inquiries/2017-04-04-Inquiry-into-MSD-collection-of-individual-client-level-data-OPC-report.pdf>.
- Privacy International. (2020, March 3). One Year On, what has Uganda’s Data Protection Law Changed? Retrieved 14 April 2022 from <https://www.privacyinternational.org/news-analysis/3385/one-year-what-has-ugandas-data-protection-law-changed>.
- Rajput, A., Li, Q., Ahvanooy, M., & Masood, I. (2019). EACMS: Emergency Access Control Management System for Personal Health Record Based on Blockchain. *IEEE Access*, 7, 84304–84317.
- Rashbrooke, M. (2021, February 8). How Taiwan is inoculating itself against the Uber “virus”. Retrieved 14 April 2022 from <https://citymonitor.ai/economy/how-taiwan-inoculating-itself-against-uber-virus-2786>.
- Raval, S. (2016). *Decentralized Applications : Harnessing Bitcoin’s Blockchain Technology*. Sebastopol, CA, USA: O’Reilly.
- Raymond, E. (2001). *The Cathedral and the Bazaar: Musings on Linux and Open Source by an accidental revolutionary* (Revised ed.). Sebastopol, CA, USA: O’Reilly.
- Reader, R. (2021, 19 February). With a wearable reportedly in the works, Facebook continues a quiet push into health. *Fast Company*. Retrieved 11 May 2022 from <https://www>

References

- [.fastcompany.com/90606044/facebook-health-tracker-wearable](https://www.fastcompany.com/90606044/facebook-health-tracker-wearable).
- Reeves, J., Treharne, G. J., Theodore, R., Edwards, W., Ratima, M., & Poulton, R. (2022). Understanding the data-sharing debate in the context of Aotearoa/New Zealand: a narrative review on the perspectives of funders, publishers/journals, researchers, participants and Māori collectives. *Kotuitui: New Zealand Journal of Social Sciences Online*, 17(1), 1–23.
- Republic of Estonia Information System Authority. (2019, July). Introduction of X-tee. Retrieved 11 May 2022 from <https://www.ria.ee/en/state-information-system/x-tee/introduction-x-tee.html>.
- Republic of Estonia Information System Authority. (2022, March 17). Data Exchange Layer X-tee. Retrieved 26 April 2022 from <https://www.ria.ee/en/state-information-system/x-tee.html>.
- Ribak, R., & Turow, J. (2003). Internet Power and Social Context: A Globalization Approach to Web Privacy Concerns. *Journal of Broadcasting & Electronic Media*, 47(3), 328–349.
- Rice, N. (2019, October 10). Estonia’s Digital Embassies and the Concept of Sovereignty. *Georgetown Security Studies Review*. Retrieved 14 April 2022 from <https://georgetownsecuritystudiesreview.org/2019/10/10/estonias-digital-embassies-and-the-concept-of-sovereignty/>.
- Richards, L. (2009). *Handling qualitative data: A practical guide* (2nd ed.). Thousand Oaks, CA, USA: Sage.
- Rivest, R., Shamir, A., & Adleman, L. (1978). A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *Communications of the ACM*, 21(2), 120–126.
- Robson, C., & McCartan, K. (2016). *Real world research: A resource for users of social research methods in applied settings* (4th ed.). Chichester, UK: Wiley.
- Rosato, D. (2020, January 28). What Your Period Tracker App Knows About You. *Consumer Reports*. Retrieved 11 May 2022 from <https://www.consumerreports.org/health-privacy/what-your-period-tracker-app-knows-about-you>.
- Roudik, P. (2016, March 2). ECHR, Russian Federation: Breaches of Human Rights in Surveillance Legislation. *The Library of Congress Global Legal Monitor*. Retrieved 11 May 2022 from <https://www.loc.gov/law/foreign-news/article/echr-russian-federation-breaches-of-human-rights-in-surveillance-legislation/>.
- Royden, L., & Li, M. (2017). Extreme Maps. *Brennan Center for Justice at New York University School of Law*. Retrieved 11 May 2022 from https://www.brennancenter.org/sites/default/files/2019-08/Report_Extreme%20Maps%205.16_0.pdf.
- Rubin, H., & Rubin, I. (2012). *Qualitative interviewing: The art of hearing data* (3rd ed.). Thousand Oaks, CA, USA: Sage.

References

- Saldana, J. (2015). *The Coding Manual for Qualitative Researchers* (3rd ed.). Thousand Oaks, CA, USA: Sage.
- Samory, M., Sen, I., Kohne, J., Flöck, F., & Wagner, C. (2021, June 7-10). “Call me sexist, but...”: Revisiting Sexism Detection Using Psychological Scales and Adversarial Samples. In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media (ICWSM2021)* (pp. 573–584).
- Sanderson, S., Brothers, K., Mercaldo, N., Clayton, E., Antommara, A., Aufox, S., ... Holm, I. (2017). Public Attitudes toward Consent and Data Sharing in Biobank Research: A Large Multi-site Experimental Survey in the US. *The American Journal of Human Genetics*, 100, 414–427. Retrieved from <http://dx.doi.org/10.1016/j.ajhg.2017.01.021>
- Sangers, T., Reeder, S., van der Vet, S., Jhingoer, S., antien Mooyaart, Siegel, D., ... Wakke, M. (2022). Validation of a Market-Approved Artificial Intelligence Mobile Health App for Skin Cancer Screening: A Prospective Multicenter Diagnostic Accuracy Study. *Dermatology*. doi: 10.1159/000520474
- Sankoh, O., & Ijsselmuiden, C. (2011, July). Sharing research data to improve public health: A perspective from the global south. *The Lancet*, 378(9789). Retrieved from [https://www.thelancet.com/journals/lancet/article/PIIS0140-6736\(11\)61211-7/fulltext](https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(11)61211-7/fulltext)
- Saputro, R., Pappel, I., Vainsalu, H., Lips, S., Draheim, D., et al. (2020, April 22-24). Prerequisites for the Adoption of the X-Road Interoperability and Data Exchange Framework: A Comparative Study. In *Proceedings of the 2020 Seventh International Conference on eDemocracy & eGovernment (ICEDEG)* (pp. 216–222). Buenos Aires, Argentina.
- Satariano, A., Mozur, P., & Krolik, A. (2022, 28 March). When Nokia pulled out of Russia, a vast surveillance system remained. *The New York Times*. Retrieved 14 April 2022 from <https://www.nytimes.com/2022/03/28/technology/nokia-russia-surveillance-system-sorm.html>.
- Savelyev, A. (2016). Russia’s new personal data localization regulations: A step forward or a self-imposed sanction? *Computer Law and Security Review*, 32, 128–145.
- Schia, N. N., & Gjesvik, L. (2017). China’s cyber sovereignty. *Norwegian Institute of International Affairs. Policy Brief [2/2017]*. Retrieved from <https://www.jstor.org/stable/resrep07952>.
- Schiffer, Z. (2021, 13 January). Period tracking app settles charges it lied to users about privacy. *The Verge*. Retrieved 11 May 2022 from <https://www.theverge.com/2021/1/13/22229303/flo-period-tracking-app-privacy-health-data-facebook-google>.
- Schork, S., & Kirchner, E. (2018, August 14-17). Defining Requirements in Prototyping: The Holistic Prototype and Process Development. In *Proceedings of NordDesign 2018*.

References

- Linköping, Sweden.
- Schultz, J., & Rainie, S. C. (2014). The strategic power of data: A key aspect of sovereignty. *The International Indigenous Policy Journal*, 5(4), 1–3.
- Schweik, C. (2011). Free/Open-Source Software as a framework for establishing commons in science. In C. Hess & E. Ostrom (Eds.), *Understanding knowledge as a commons: From theory to practice* (pp. 277–310). Cambridge, MA, USA: MIT Press.
- Scoon, C., & Ko, R. (2016, August). The Data Privacy Matrix Project: Towards a Global Alignment of Data Privacy Laws. *2016 IEEE Trustcom/BigDataSE/ISPA*, 1998–2005.
- Seedat, S., & Rondon, M. (2021). Women’s wellbeing and the burden of unpaid work. *British Medical Journal*, 374. Retrieved 11 May 2022 from <https://www.bmj.com/content/374/bmj.n1972>.
- Selke, M. (2018, February 7). Here’s why the Estonian Police and Border Guard Board now has a presence in South Korea. *Republic of Estonia E-Residency*. Retrieved 11 May 2022 from <https://medium.com/e-residency-blog/heres-why-the-estonian-police-and-border-guard-board-now-has-a-presence-in-south-korea-ead9b3dc8387>.
- Sethi, A. (2011, December 22). The False Promise of Biometrics. *The New York Times*. Retrieved 21 April 2019 from <https://latitude.blogs.nytimes.com/2011/12/22/the-false-promise-of-biometrics-in-india/>.
- Sharma, V., & Mukherjee, S. (2018, June). GitLab gains developers after Microsoft buys rival GitHub. *Reuters*. Retrieved 11 May 2022 from <https://www.reuters.com/article/us-github-microsoft-gitlab-idUSKCN1J12BR>.
- Simon, H. (1996). *The sciences of the artificial* (3rd ed.). Cambridge, MA, USA: MIT Press.
- Social Investment Agency. (2018, November). What you told us. *Social Investment Agency*. Retrieved 21 April 2021 from <https://swa.govt.nz/assets/Uploads/what-you-told-us.pdf>.
- Social Wellbeing Agency. (2022, January). Data Protection & Use Policy. Retrieved 26 April 2022 from <https://www.digital.govt.nz/assets/Standards-guidance/Privacy/Data-Protection-and-Use-Policy-DPUP-January-2022-Version-1.2.pdf>.
- Soldatov, A., & Borogan, I. (2013). Russia’s surveillance state. *World Policy Journal*, 30(3). Retrieved from <https://worldpolicy.org/2013/09/12/russias-surveillance-state/>.
- Solon, O. (2018, July 13). ‘Data is a fingerprint’: why you aren’t as anonymous as you think online. *The Guardian*. Retrieved 11 May 2022 from <https://www.theguardian.com/world/2018/jul/13/anonymous-browsing-data-medical-records-identity-privacy>.

References

- Sporle, A., Hudson, M., & West, K. (2021). Indigenous data and policy in Aotearoa New Zealand. In M. Walter, T. Kukutai, S. R. Carroll, & D. Rodriguez-Lonebear (Eds.), *Indigenous Data Sovereignty and Policy* (pp. 62–80). Abingdon, Oxon, UK: Routledge.
- Statistics New Zealand. (2018, August). A social licence approach to trust. Retrieved 11 May 2022 from <https://www.stats.govt.nz/assets/Uploads/Corporate/Measuring-Stats-NZs-social-licence/a-social-licence-approach-to-trust.pdf>.
- Statistics New Zealand. (2020, April 30). 2018 Census totals by topic – national highlights. Retrieved 11 May 2022 from <https://www.stats.govt.nz/information-releases/2018-census-totals-by-topic-national-highlights-updated>.
- Strongman, S. (2018, 23 February). Why this year's census won't ask questions about sexual orientation, diverse gender and sexual identities. *Radio New Zealand*. Retrieved 11 May 2022 from <https://www.rnz.co.nz/news/the-wireless/375135/why-this-year-s-census-won-t-ask-questions-about-sexual-orientation-diverse-gender-and-sexual-identities>.
- Swan, J., & Newell, S. (1996). The role of professional associations in technology diffusion. *Organization Studies*, 16(5), 847–874.
- Sweeney, L. (2015). Only You, Your Doctor, and Many Others May Know. *Technology Science*. Retrieved 11 May 2022 from <https://techscience.org/a/2015092903>.
- Taiwan National Development Council. (2022, May 5). MyData. Retrieved 7 May 2022 from https://www.ndc.gov.tw/en/Content_List.aspx?n=3E2BA2908A49AC6A.
- Talmazan, Y. (2019, June 26). Data security meets diplomacy: Why Estonia is storing its data in Luxembourg. *NBC News*. Retrieved 16 April 2022 from <https://www.nbcnews.com/news/world/data-security-meets-diplomacy-why-estonia-storing-its-data-luxembourg-n1018171>.
- Tan, A. (2015, February 14). Why Some Native Americans Say Facebook Is Biased Against Them. *ABC News*. Retrieved 7 May 2022 from <https://abcnews.go.com/Technology/native-americans-petition-facebook-cease-deactivations-names/story?id=28921793>.
- Tan, Z., & Zhang, C. (2021). China's PIPL and DSL: Is China following the EU's approach to data protection? *Journal of Data Protection & Privacy*, 5(1), 7–25.
- Tang, A. (2019, March 22). Inside Taiwan's new digital democracy. *The Economist*. Retrieved 7 May 2022 from <https://www.economist.com/open-future/2019/03/12/inside-taiwans-new-digital-democracy>.
- Tarr, D., Lavoie, E., Meyer, A., Tschudin, C., et al. (2019, September 24–26). Secure ScuttleButt: An identity-centric protocol for subjective and decentralized applications. In *Proceedings of the 2019 Conference on Information-Centric Networking, ICN 2019* (pp. 1–11). Macao,

References

- China.
- Taylor, L., & Broeders, D. (2015). In the name of Development: Power, profit and the datafication of the global South. *Geoforum*, 64, 229–237.
- Terhune, C., Epstein, K., & Arnst, C. (2009, May 4). The dubious promise of digital medicine. *Business Week*, 30–37.
- The European Institute for Innovation through Health Data. (2021). The European Institute for Innovation through Health Data home page. Retrieved from <https://www.i-hd.eu/>.
- The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. (2019). *Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent system*. Retrieved 9 April 2022 from <https://standards.ieee.org/wp-content/uploads/import/documents/other/eadle.pdf>.
- The Office for Māori Crown Relations - Te Arawhiti. (2018). *Guidelines for engagement with Māori* (Tech. Rep.). Retrieved 11 May 2022 from <https://www.tearawhiti.govt.nz/assets/Maori-Crown-Relations-Roopu/6b46d994f8/Engagement-Guidelines-1-Oct-18.pdf>.
- Theiner, J., Müller-Budack, E., & Ewerth, R. (2022, January 3-8). Interpretable Semantic Photo Geolocation. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (pp. 1474–1484). Hawaii, USA. doi: 10.1109/WACV51458.2022.00154
- Thompson, M., & Lockwood, P. (2021, 12 April). China hits Alibaba with record \$2.8 billion fine for behaving like a monopoly. *CNN*. Retrieved 7 April 2022 from <https://edition.cnn.com/2021/04/10/tech/alibaba-china-record-fine/index.html>.
- Thuan, N. H., Drechsler, A., & Antunes, P. (2019). Construction of Design Science Research Questions. *Communications of the Association for Information Systems*, 44, 332–363.
- Tikk, E. (2021). The leaps and bounds of e-Estonia. *Media Development*, 67(2), 29–32.
- Toh, M. (2016, November 14). A Practical Guide to the Singapore Personal Data Protection Act (PDPA). *Asia Law Network*. Retrieved 11 May 2022 from <http://learn.asialawnetwork.com/2016/11/14/definite-guide-singapore-pdpa-personal-data-protection-act/>.
- Turow, J., Feldman, L., & Meltzer, K. (2005, June 1). Open to Exploitation: America's Shoppers Online and Offline. *A Report from the Annenberg Public Policy Center of the University of Pennsylvania*. Retrieved 11 May 2022 from https://repository.upenn.edu/asc_papers/35.
- Turow, J., Hennessy, M., & Draper, N. (2015). The Tradeoff Fallacy: How Marketers are Misrepresenting American Consumers and Opening Them Up to Exploitation. *A Report from the Annenberg School for Communication, University of Pennsylvania*. Retrieved 11 May 2022 from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2820060.

References

- Unique Identification Authority of India. (2022, April 26). Aadhaar dashboard. Retrieved on April 26 2022 from https://uidai.gov.in/aadhaar_dashboard/.
- United Nations Office of the High Commissioner for Human Rights. (2018). A human-rights based approach to data. Retrieved 11 May 2022 from <https://www.ohchr.org/sites/default/files/Documents/Issues/HRIndicators/GuidanceNoteonApproachtoData.pdf>.
- Unwanted Witness. (2021, November). Privacy Scorecard Report. Retrieved 14 April 2022 from <https://www.unwantedwitness.org/wp-content/uploads/2021/11/Privacy-Scorecard-Report-2021.pdf>.
- Vassil, K. (2015, June). Estonian e-Government Ecosystem: Foundation, Applications, Outcomes. *World Bank*. Retrieved 11 May 2022 from <http://pubdocs.worldbank.org/en/165711456838073531/WDR16-BP-Estonian-eGov-ecosystem-Vassil.pdf>.
- Vaughan-Nicholls, S. (2018, February). Open source is 20: How it changed programming and business forever. *ZDNet*. Retrieved from <https://www.zdnet.com/article/open-source-turns-20>
- Venable, J., Pries-Heje, J., & Baskerville, R. (2012, 05). A Comprehensive Framework for Evaluation in Design Science Research. In *Proceedings of the 7th International DESRIST Conference on Design Science Research in Information Systems: Advances in Theory and Practice (DESRIST 2012)* (Vol. 7286, pp. 423–438). doi: 10.1007/978-3-642-29863-9_31
- Vogelstein, F. (2018, December 19). Why Should Anyone Believe Facebook Anymore? *Wired*. Retrieved 11 May 2022 from <https://www.wired.com/story/facebook-data-sharing-privacy-investigation>.
- W3Techs. (2022, 7 May). Usage statistics of web servers. Retrieved 7 May 2022 from https://w3techs.com/technologies/overview/web_server.
- Wagner, J. (2017, 1 June). China's Cybersecurity Law: What You Need to Know. *The Diplomat*. Retrieved 18 April from <https://thediplomat.com/2017/06/chinas-cybersecurity-law-what-you-need-to-know/>.
- Wahlstrom, K., Ulhaq, A., & Burmeister, O. (2020). Privacy by design: a Holochain exploration. *Australasian Journal of Information Systems*, 24, 1–9.
- Waitangi Tribunal. (2019). Hauora: Report on Stage One of the Health Services and Outcomes Kaupapa Inquiry. *Waitangi Tribunal*. Retrieved 26 April 2022 from https://forms.justice.govt.nz/search/Documents/WT/wt_DOC_152801817/Hauora%20W.pdf.
- Walch, A. (2019). Deconstructing “Decentralization”. In C. Brummer (Ed.), *Cryptoassets: Legal, Regulatory, and Monetary Perspectives* (pp. 39–68). Oxford Scholarship Online.
- Walter, M., & Carroll, S. R. (2021). Indigenous Data Sovereignty, governance and the link to In-

References

- digenous policy. In M. Walter, T. Kukutai, S. R. Carroll, & D. Rodriguez-Lonebear (Eds.), *Indigenous Data Sovereignty and Policy* (pp. 1–20). Abingdon, Oxon, UK: Routledge.
- Wang, H., Zheng, Z., Xie, S., Dai, H.-N., & Chen, X. (2018, 10). Blockchain challenges and opportunities: a survey. *International Journal of Web and Grid Services*, 14, 352–375. doi: 10.1504/IJWGS.2018.10016848
- Wang, Z., Hale, S. A., Adelani, D., Grabowicz, P. A., Hartmann, T., Flöck, F., & Jurgens, D. (2019, May 13-17). Demographic Inference and Representative Population Estimates from Multilingual Social Media Data. In *Proceedings of the 2019 World Wide Web Conference (WWW '19)* (pp. 2056–2067). San Francisco, USA.
- Waring, M. (2018). Still counting : wellbeing, women’s work and policy-making.
- Warren, S., & Brandeis, L. (1890). The Right to Privacy. *Harvard Law Review*, 4(5), 193–220.
- Warren, T. (2020, May). Microsoft: we were wrong about open source. *The Verge*. Retrieved 11 May 2022 from <https://www.theverge.com/2020/5/18/21262103/microsoft-open-source-linux-history-wrong-statement>.
- Weber, R. (2010). New sovereignty concepts in the age of the internet? *Journal of Internet Law*, 14(8), 12–20.
- West, P., Van Kleek, M., Giordano, R., Weal, M., Shadbolt, N., et al. (2017). Information Quality Challenges of Patient-Generated Data in Clinical Practice. *Frontiers in Public Health*, 5. Retrieved from <https://www.frontiersin.org/article/10.3389/fpubh.2017.00284> doi: 10.3389/fpubh.2017.00284
- Whisnant, R. (2010). "A Woman’s Body Is Like a Foreign Country" Thinking about National and Bodily Sovereignty. In P. DesAuteles & R. Whisnant (Eds.), *Global Feminist Ethics* (p. 155-176). Rowman & Littlefield.
- Williams, J., Vis-Dunbar, M., & Weber, J. (2011). First nations privacy and modern health care delivery. *Indigenous Law Journal*, 10(1), 101–132.
- Winter, R. (2008). Design Science Research in Europe. *European Journal of Information Systems*, 17(5), 470–475. doi: 10.1057/ejis.2008.44
- Wolf, J. A., Moreau, J. F., Akilov, O., Patton, T., English, I., Joseph C., Ho, J., & Ferris, L. K. (2013, 04). Diagnostic Inaccuracy of Smartphone Applications for Melanoma Detection. *JAMA Dermatology*, 149(4), 422–426. Retrieved from <https://doi.org/10.1001/jamadermatol.2013.2382> doi: 10.1001/jamadermatol.2013.2382
- Wong, K. O., Zaïane, O. R., Davis, F. G., & Yasui, Y. (2020). A machine learning approach to predict ethnicity using personal name and census location in Canada. *PloS one*, 15(11), 1–16.
- Woods, S., Schwartz, E., Tuepker, A., Press, N., Nazi, K., Turvey, C., & Nichol, W. (2013). Patient Experiences With Full Electronic Access to Health Records and Clinical Notes

References

- Through the My HealtheVet Personal Health Record Pilot: Qualitative Study. *Journal of Medical Internet Research*, 15(3). doi: 10.2196/jmir.2356
- World Economic Forum. (2016, September). Values and the Fourth Industrial Revolution: Connecting the Dots Between Value, Values, Profit and Purpose. Retrieved 11 May 2022 from https://www3.weforum.org/docs/WEF_Values_and_the_Fourth_Industrial_Revolution_WHITEPAPER.pdf.
- Xia, Q., Sifah, E. B., Asamoah, K. O., Gao, J., Du, X., & Guizani, M. (2017). MeDShare: Trust-Less Medical Data Sharing Among Cloud Service Providers via Blockchain. *IEEE Access*, 5, 14757–14767. doi: 10.1109/ACCESS.2017.2730843
- Xiao, E. (2021, 17 August). China Set to Pass One of the World’s Strictest Data-Privacy Laws; New curbs come as citizens grow increasingly concerned about tech giants’ nosiness. *The Wall Street Journal*.
- xkcd. (n.d.). Automation. *xkcd*. Retrieved 11 May 2022 from <https://xkcd.com/1319/>.
- Yang, F., & Xu, J. (2018). Privacy concerns in China’s smart city campaign: The deficit of China’s Cybersecurity Law. *Asia and the Pacific Policy Studies*, 5(3), 533–543.
- Yegulalp, S. (2018, November 28). Rust language is too hard to learn and use, says user survey. *InfoWorld*. Retrieved 14 April 2022 from <https://www.infoworld.com/article/3324488/rust-language-is-too-hard-to-learn-and-use-says-user-survey.html>.
- Yliaska, V. (2015). New Public Management as a Response to the Crisis of the 1970s: The Case of Finland, 1970–1990. *Contemporary European History*, 24(3), 435–459.
- Yurieff, K. (2018, September 20). Google still lets third-party apps scan your Gmail data. *CNN*. Retrieved 11 May 2022 from <https://money.cnn.com/2018/09/20/technology/google-gmail-scanning/index.html>.
- Zeng, E., Mare, S., & Roesner, F. (2017, July). *End User Security and Privacy Concerns with Smart Homes*. Paper presented at the Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017). Santa Clara, USA. Retrieved from <https://www.usenix.org/system/files/conference/soups2017/soups2017-zeng.pdf>
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York, NY: PublicAffairs.

A. Survey Participant Information Sheet



Participant Information Sheet

Date Information Sheet Produced:

23 August 2020

Project Title

Designing, building and implementing a Radically Decentralised Health Information System in New Zealand

An Invitation

My name is Alex Poor, and I currently work within Pinnacle as the Chief Data Officer.

At the same time, I am studying for a PhD at AUT and would really value the opportunity to seek your views on my research topic.

In this PhD I am hoping to identify how we can build a data ecosystem that lets individuals take control of their own health data. I intend to build a system that would let you take your own data (from blood pressure monitors, smart scales or FitBit devices) and hold it together in a secure place – alongside their GP and hospital data – that you have full control over. You would be able to share that data with whomever you choose, and you would also be able to see exactly how it is being used if you do share it.

To help design this system, I am seeking your views on whether or not you think this would be useful – and what should be considered in such a system.

What is the purpose of this research?

The results of this survey will be analysed, and will inform the design of the prototype system. The prototype will then be evaluated separately.

This research is directly related to my PhD study, and will form a core part of the final thesis. The results may also be used as part of academic publications, or presentations.

How was I identified and why am I being invited to participate in this research?

As a Pinnacle staff member I am hoping to access the views and experience of the 'expert patients' on this Consumer Panel. Permission has been sought from Pinnacle to offer my survey to the Panel to fulfil the research requirements noted above.

How do I agree to participate in this research?

Your participation in this research is completely voluntary (it is your choice) and whether or not you choose to participate will neither advantage nor disadvantage you.

If you do choose to take part, the survey itself will ask you to reconfirm your consent.

Since this is anonymous survey, any data you submit cannot be retrieved, returned to you or removed since there will be no way to identify which response belongs to you. Therefore, please be certain about your participation in the survey before you provide consent and complete the questionnaire.

What will happen in this research?

This research invites you to complete a brief electronic survey. The link to access the survey will be sent to you, via Pinnacle.

If you consent to take part, please follow the provided link and complete the survey to the best of your ability.

A. Survey Participant Information Sheet

What are the discomforts and risks?

There are unlikely to be any discomforts, or risks, particularly if you are familiar with how the Consumer Panel surveys ordinarily run.

How will these discomforts and risks be alleviated?

If you do experience any discomfort then please refrain from accessing or completing the survey and, additionally, speak with myself or my supervisor (details below).

What are the benefits?

In this research I want to highlight an important issue that is not well researched. Currently, there are initiatives to give people better access to their health data (for example, via patient portals such as ManageMyHealth or MyIndici). However, these tools do not give real control to people – there is a lot of research indicating that, if people have full control of their data, they feel more empowered and engaged with staying well. There are no current initiatives or research projects focusing on giving individuals this level of control.

The research therefore fills an important gap, as we move further and further into using technology to improve healthcare – all fuelled by the capture of data.

The benefit to myself personally is that my research will lead to graduating as a PhD, and the findings may be presented in journal articles or at conferences.

How will my privacy be protected?

As the researcher, I have no access to the Consumer Panel information held by Pinnacle. This survey will be completely anonymous – you are not required or invited to enter an email address or any other identifying information.

You will notice there are some brief questions asking for information such as age group, gender and ethnicity. These are not mandatory, but they do help my examiners determine if the survey is 'representative', and the results can be extrapolated across New Zealand. If you have any concerns about providing this information, please feel free to leave these questions blank.

What are the costs of participating in this research?

The survey will take approximately thirty minutes to complete.

What opportunity do I have to consider this invitation?

The survey will have an expiry date, which will be three weeks from the date it is sent by Pinnacle. After this time it will not be possible to access it.

Will I receive feedback on the results of this research?

Yes.

A brief summary of the survey results will be displayed on a public web page, once analysis has been completed. A link to this web page will be provided to Pinnacle, who will forward it to the Consumer Panel.

What do I do if I have concerns about this research?

Any concerns regarding the nature of this project should be notified in the first instance to the Project Supervisor, Professor Marilyn Waring, marilyn.waring@aut.ac.nz, (+649) 921 9999 ext 9661.

Concerns regarding the conduct of the research should be notified to the Executive Secretary of AUTC, ethics@aut.ac.nz, (+649) 921 9999 ext 6038.

A. Survey Participant Information Sheet

Whom do I contact for further information about this research?

Please keep this Information Sheet for your future reference. You are also able to contact the research team as follows:

Researcher Contact Details:

Alex Poor, fy5275@autuni.ac.nz, (+64 27) 297 4591.

Project Supervisor Contact Details:

Professor Marilyn Waring, marilyn.waring@aut.ac.nz, (+649) 921 9999 ext 9661.

Approved by the Auckland University of Technology Ethics Committee on 10 July 2020, AUTEC Reference number 20/201.

B. Survey: EA1 Approval



Auckland University of Technology Ethics Committee (AUTC)

Auckland University of Technology
D-88, Private Bag 92006, Auckland 1142, NZ
T: +64 9 921 9999 ext. 8316
E: ethics@aut.ac.nz
www.aut.ac.nz/researchethics

10 July 2020

Marilyn Waring
Faculty of Culture and Society

Dear Marilyn

Ethics Application: 20/201 Designing, building and implementing a Radically Decentralised Health Information System in New Zealand

Thank you for submitting your application for ethical review. We are pleased to advise that a subcommittee of the Auckland University of Technology Ethics Committee (AUTC) approved your ethics application in stages, subject to the following conditions:

1. Provision of an assurance that post-analysis data will be stored securely on an external hard drive on AUT premises;
2. Inclusion of the Information Sheet as the first page of the survey. This can be an abbreviated version but must contain the following components:
 - a. an overview of the research and what it is you wish them to do;
 - b. the AUT logo;
 - c. the section on concerns and the AUTC approval details, wording for which can be found in the Information Sheet template on the Research Ethics website at <http://aut.ac.nz/researchethics>;
 - d. a statement noting that participants may withdraw from the research at any point until they submit their responses but that once this has been done their data can neither be identified nor withdrawn.

This approval is for the survey phase of the research. Full information about future stages of this research needs to be provided to and approved by AUTC before the data collection for those stages commences.

Please provide us with a response to the points raised in these conditions, indicating either how you have satisfied these points or proposing an alternative approach. AUTC also requires copies of any altered documents, such as Information Sheets, surveys etc. You are not required to resubmit the application form again. Any changes to responses in the form required by the committee in their conditions may be included in a supporting memorandum.

Please note that the Committee is always willing to discuss with applicants the points that have been made. There may be information that has not been made available to the Committee, or aspects of the research may not have been fully understood.

Once your response is received and confirmed as satisfying the Committee's points, you will be notified of the full approval of your ethics application. Full approval is not effective until all the conditions have been met. Data collection may not commence until full approval has been confirmed. If these conditions are not met within six months, your application may be closed and a new application will be required if you wish to continue with this research.

To enable us to provide you with efficient service, we ask that you use the application number and study title in all correspondence with us. If you have any enquiries about this application, or anything else, please do contact us at ethics@aut.ac.nz.

We look forward to hearing from you,

B. Survey: EA1 Approval

(This is a computer-generated letter for which no signature is required)

The AUTC Secretariat

Auckland University of Technology Ethics Committee

Cc: alex@robot5x.com; Rhema Vaithianathan; Dave Parry

C. Survey final report

Draft survey results – as at 8/10/2020

Background

The Consumer Panel is **1,671** individuals in total – all were invited to take part. **426** did, giving a response rate of **25%**.

Ethnicity – very non-representative!

Ethnicity	Count	%
asian	1	0.2%
euro	389	91.3%
maori	6	1.4%
mela	3	0.7%
other	18	4.2%
pacific	3	0.7%
(empty)	6	1.4%
Total Result	426	100.0%

Please note the Consumer Panel overall is very skewed towards Europeans. A total of 131 (7.8%) Panel members identify as Maori versus only 6 (1.4%) identifying as Maori in my survey. Please note, however, that the Consumer Panel data allows selection of multiple ethnicities whereas my survey only permitted one.

Gender – majority female

Gender	Count	%
female	268	62.9%
male	153	35.9%
decline	1	0.2%
(empty)	4	0.9%
Total Result	426	100.0%

Age group – skewed towards older people

Age group	Count	%
15-24	3	0.7%
25-44	30	7.0%
45-64	144	33.8%
65+	228	53.5%
(empty)	21	4.9%
Total Result	426	100.0%

C. Survey final report

Top 5 subgroups:

28.4% - European females aged 65+
21.4% - European males aged 65+
20.9% - European females aged 45-64
8.0% - European males aged 45-64
5.9% - European females aged 25-44

Question responses

"I want to be able to own my health data, and decide who can access it"

Overall **71.1%** of respondents agreed or strongly agreed with this statement.
There was no significant variation when broken down by age group (with 72.7% of those aged 15-44 agreeing or strongly agreeing).

By ethnicity, all groups except European and Pacific tended to agree more strongly with this statement – however the actual numbers are very small (eg. All 3 MELAA respondents agreed, and 88.9% of the 18 'Other' respondents agreed).

13.6% felt neutrally about this overall, while **12.4%** disagreed or disagreed strongly.

"I want to be able to record and manage my own health data (eg. from smart devices or wearables) and see it alongside all my other medical information"

Overall **61%** of respondents agreed or strongly agreed with this statement.
There was no variation when broken down by age group, with all subcategories within 3% of the overall result.

By ethnicity, the 'Other' group agreed more often with this statement (**72.2%**).
By gender there was no significant difference, however females left this question empty more often than males (**5.2%** vs **1.3%**).

20.4% felt neutrally about this overall, while **14.8%** disagreed or disagreed strongly.

C. Survey final report

"I want to be able to see who is using my health data"

78.2% agreed or agreed strongly with this statement.

By age group, 65+ agreed with this statement the most (at **82%**), perhaps a product of using the health system more and having more opportunity to feel frustration around health information. By contrast **24.2%** of the 15-44 age group either disagreed or felt neutral about this.

By ethnicity, the 'Other' group agreed more often with this statement (**88.9%**). Although a very small group, 2 of the 3 Pacific respondents disagreed *strongly* with this statement.

By gender, females tended to agree more than males (**80.2%** vs **75.2%**), whilst males had a relatively high proportion of neutral or empty responses (**14.4%**).

The next set of questions asks respondents to imagine that they can control access to their health data. They are asked to indicate how (or if) they would share data with a range of different entities.

"How would you share your data with your GP?"

89% said they would give full and ongoing access to their GP, and this is positively correlated with age group. For example, **80%** of 25-44 year olds going up to **90.8%** of 65+ year olds.

By ethnicity there is no significant variation however, of those not supplying an ethnicity, only **66.7%** (out of a total of 6) agreed to full and ongoing GP access.

There is virtually no variation at all by gender.

"How would you share your data with a hospital team?"

There was slightly less support for this, with **59.4%** offering to provide full and ongoing access, but **35%** saying they would be happy to share their data *only when required*.

3 of 8

C. Survey final report

There is some variation by age group. 25-44 year olds were least likely to support ongoing access to a hospital team (**46.7%**), whereas 15-24 year olds were the most likely (**66.7%**). This question also sees a large gap between the 45-64 and the 65+ age group, with **51.4%** support for ongoing access in the former and a much higher **64.9%** in the latter.

By gender there is a key difference whereby males are more supportive of ongoing access (at **66%**) versus **55.2%** for females. This is driven by males in the 65+ category (who, it should be remembered, also make up the largest number of male respondents).

“How would you share your data with your whanau or family?”

48.6% of respondents said they'd give full and ongoing access to their health data on an ongoing basis. Additionally **34.3%** would give access whenever required.

Again this seems to be positively correlated with age, where only **30%** of 25-44 year olds would give ongoing access to their family, compared with **54.8%** of 65+ year olds. This 25-44 group also had the highest rate of answering 'Never' to this question (**6.7%**).

There is also a key difference between genders, whereby males are again far more supportive of ongoing access (at **64.1%**) versus **39.9%** for females. The support of males in this area is concentrated in older age groups, with **71.7%** of males age 65+ supportive of ongoing access.

“How would you share your data with the Ministry of Health about your hospital visits?”

This question is intended to elicit attitudes around population of the Ministry's national data collections. Currently this data is collected automatically by DHB's and forwarded monthly to the centre for certain data sets.

This item generated more diversity of opinion, with only **30.3%** saying they would provide ongoing access and **4.7%** answering 'Never'.

Countering some of the other correlations with age, younger age groups would give ongoing access more often than older age groups (**42.4%** of 15-44 year olds, versus **29.8%** of 45+ year olds). The group answering 'Never' to this question the most were aged 45-64 (**7.6%**).

C. Survey final report

Once again we see males being more supportive of ongoing access (at **36%**) versus **26.5%** for females.

“How would you share your data with other government agencies (eg. MSD or ACC)?”

This question provoked the strongest negative response, with **15.5%** answering ‘Never’, and only **12.4%** offering full and ongoing access.

Younger age groups were somewhat more open to this concept, with **27.3%** of 15-44 year olds open to ongoing access, compared with **9%** for those aged 45-64.

Contrary to some of the previous questions, there was negligible difference in response when broken down by gender.

“How would you share your data anonymously to a government database, so it can be combined with other data and used for policy analysis and research...”

This question is intended to elicit attitudes around the Integrated Data Infrastructure (IDI). Currently this aggregates official data sets and deidentifies personal information so it can be used for research and policy analysis.

I must admit I anticipated a broadly negative response to this, but there was a very solid **38.3%** who stated they would provide full and ongoing access to their data for this purpose. Conversely, a relatively high **14.6%** said they would ‘Never’ share their data in this way.

This pattern is magnified by age, where younger age groups were both slightly more happy to share their data (**40%** of 25-44 year olds would share on an ongoing basis) and also considerably more opposed (**30%** of 25-44 year olds answered ‘Never’).

Males are again happier to share their data in this domain, with **45.8%** of males supportive of ongoing access versus **33.6%** of females.

C. Survey final report

Interesting comments

ACC are the most underhand people in the government system.

While I am happy for my husband to have full access to my data at the moment, I don't necessarily want other members of my family to have total access. There might also come a time in the future when I don't want my husband to have full access.

i don't completely trust in govt or NGO organisations in their ability to keep public data safe from accidental release or abuse by 3rd parties.

I guess if I had a chronic illness I would want MSD or ACC to have access so that I did not have to reapply for benefits - they could just check the details from my health records.

My data to be accessed by govt would need to be upon request only, each time reqd., with right to deny access at each given time

We have seen how Privacy can be lost through human error. ACC could easily use data for their own purposes to deny claims etc

I simply do not trust the government to protect my privacy, there have been far too many instances where privacy has been breached, and individuals I for.ation has been made public.

I don't have a problem with sharing my data with relevant people or organisations. I need to know they will not share it with ANYONE else unless I give my approval.

Government have a bad record at keeping information secure. They should never be allowed to access my data without my express permission

I want to know what the government and agencies are doing with my data before they have it.

It's mine, it's about me, I should own it absolutely and outright.

Our present system of our own medical data has been very inefficient - put together by nerds I suspect, not practitioners with life experience (at almost anything). It also has virtually no data but a continuous cost.

Other than health care professionals, no one should be able to access my records unless I give explicit consent

I do not understand personal health data as something exclusive to one, it's also in part a part of a wider body and a whole system. As one who is anxiety prone, imaging me reading all that data over the years, thank god i did not had access. I even find the dr notes one can access supper stupid, they simply have not got the time to do that better if you ask me.

Given data breaches I do not want my personal data accessed without my permission y government departments if the stock exchange can be hacked so can health data... But the privacy act should not be a cloak for family members to advocate and help unwell relatives

My health data is mine and it is my choice who to share it with or to give access to it

6 of 8

C. Survey final report

I like the idea of having access to all my own records, partly to ensure they are correct and also to better understand my own health over time. I have found when I did get access to some GP notes that they had not correctly recorded what I said and worry how that might lead to incorrect diagnosis.

I am happy to share my data to improve services to enable improvements in the health system -I have nothing to hide.

I would love to be in control of my medical data. I am totally against multi agencies having any kind of access.

I may want to give my lawyer access to my health data too

I really like the idea of accessing and controlling my own data and who can see it - I do wonder if for some, the effort and decision making could put people off in practice, as they are so used to being spoon fed by applications and products created by companies like Google and Apple, that having control could be seen as "tiring" and "effort" even though it is really important. However the system is designed, it would need to be really easy for users to understand and access. Sharing with whanau would be revolutionary for how chronic conditions are managed.

I come back to It is MY health data....not theirs, not yours, not my family's...it is MINE: and I have the right to state who, how, why and when. If I cannot do that - they then do NOT have access.

C. Survey final report

Broad conclusions

People are generally happy to share all their data with their GP.

They would like to have more control over sharing data with the wider health system, and their family.

There is moderate support for sharing data with government in a deidentified fashion.

The lowest level of support was for sharing identifiable health data with other government agencies (specifically MSD or ACC).

People really want to exert control over their health data, and see who has accessed it.

People are less interested in integrating health data from different sources (eg. Wearables).

All the above conclusions based on an overwhelmingly older and European response group.

D. Interviews: EA1 Approval



Auckland University of Technology Ethics Committee (AUTEC)

Auckland University of Technology
D-88, Private Bag 92006, Auckland 1142, NZ
T: +64 9 921 9999 ext. 8316
E: ethics@aut.ac.nz
www.aut.ac.nz/researchethics

17 May 2021

Marilyn Waring
Faculty of Culture and Society

Dear Marilyn

Re Ethics Application: **21/146 Designing, building and implementing a radically decentralised health information system in New Zealand**

Thank you for providing evidence as requested, which satisfies the points raised by the Auckland University of Technology Ethics Committee (AUTEC).

Your ethics application has been approved for three years until 17 May 2024.

Standard Conditions of Approval

1. The research is to be undertaken in accordance with the [Auckland University of Technology Code of Conduct for Research](#) and as approved by AUTEC in this application.
2. A progress report is due annually on the anniversary of the approval date, using the EA2 form.
3. A final report is due at the expiration of the approval period, or, upon completion of project, using the EA3 form.
4. Any amendments to the project must be approved by AUTEC prior to being implemented. Amendments can be requested using the EA2 form.
5. Any serious or unexpected adverse events must be reported to AUTEC Secretariat as a matter of priority.
6. Any unforeseen events that might affect continued ethical acceptability of the project should also be reported to the AUTEC Secretariat as a matter of priority.
7. It is your responsibility to ensure that the spelling and grammar of documents being provided to participants or external organisations is of a high standard and that all the dates on the documents are updated.

AUTEC grants ethical approval only. You are responsible for obtaining management approval for access for your research from any institution or organisation at which your research is being conducted and you need to meet all ethical, legal, public health, and locality obligations or requirements for the jurisdictions in which the research is being undertaken.

Please quote the application number and title on all future correspondence related to this project.

For any enquiries please contact ethics@aut.ac.nz. The forms mentioned above are available online through <http://www.aut.ac.nz/research/researchethics>

(This is a computer-generated letter for which no signature is required)

The AUTEC Secretariat
Auckland University of Technology Ethics Committee

Cc: alex@robot5x.com; Rhema Vaithianathan; Dave Parry

E. Interview Participant Information Sheet



Participant Information Sheet

Date Information Sheet Produced:

28 April 2021


Project Title

Designing, building and implementing a Radically Decentralised Health Information System in New Zealand

An Invitation

My name is Alex Poor, and I am studying for a PhD at AUT; I would really value the opportunity to seek your views on my research topic.

In this PhD I am hoping to identify how I can build a data ecosystem that lets individuals take control of their own health data. I will build a prototype of a system that will let people capture their own health data (for example, from wearable devices) and hold it together in a secure place – alongside GP and Hospital data – that they have full control over. An individual would be able to share that data with whomever you choose, and would also be able to see exactly how it is being used if it is shared. This prototype will demonstrate that 'data sovereignty' is now technically possible.



Centralisation has been the dominant data model for a long time, and this has led to the embedding of policy, regulation, law, cultural practices and mental models that support and reinforce it.

The image to the left shows this centralised model in red; there is a central point which everyone has to connect to. The owner of that central point can dictate the terms of access to data, how it is stored

and how it is used. The decentralised model, shown in blue, removes this central point completely and lets people connect directly to each other whilst also being in control of how their data is accessed.

I would like to interview you, as an expert with experience and knowledge of either the New Zealand health/social sector or of decentralised technologies in general, to ask about how a decentralised system could be implemented in practice.

What is the purpose of this research?

This research is directly related to my PhD study, and will form a core part of the final thesis. The results may also be used as part of academic publications, or presentations.

These interviews will be transcribed, coded and analysed and used to develop a picture of how practical decentralisation is in the New Zealand health context.

How was I identified and why am I being invited to participate in this research?

You have been identified as someone who has particular knowledge of the New Zealand public sector data landscape, or, expert knowledge of decentralised technologies in general. I have aimed to encompass a range of perspectives in selection of participants,

E. Interview Participant Information Sheet

and tried to include those who can speak from both clinical and non-clinical viewpoints – as well as giving voice to different cultural perspectives, particularly Te Ao Māori.

How do I agree to participate in this research?

To confirm your participation in this research, please complete the Consent Form which has been emailed to you.

Your participation in this research is completely voluntary (it is your choice) and whether or not you choose to participate will neither advantage nor disadvantage you. You are able to withdraw from the study at any time. If you choose to withdraw from the study, then you will be offered the choice between having any data that is identifiable as belonging to you removed or allowing it to continue to be used. However, once the findings have been produced, removal of your data may not be possible.

What will happen in this research?

In this research I am asking you to provide approximately one hour of time in which to conduct an interview. Depending on your preferences, and your location, this can be face to face – at your place of work, or another suitable publicly accessible place – or can be done remotely using a videoconference tool (for example, Zoom). The time and place of the interview is entirely up to you. You will be sent a copy of the indicative interview questions at least one week prior to the interview itself.

The interviews are ‘semi-structured’, which means that I will have a framework of broad questions that will be posed and then each participant can answer as they choose, from their own perspective. The intention of this is not to constrain the interview, and to make the most of your own knowledge and expertise, so that it can enrich the research.

The interview audio will be recorded – if the meeting is face to face this will be done using a phone; if videoconferencing is used then computer software will be used to record the audio. This audio is then transcribed by the researcher (who may use automated transcription services as a starting point). At this point, each participant will be sent a copy of the transcript and invited to correct, redact or otherwise edit the transcription as they see fit.

The final transcript is then used by the researcher for analysis. Qualitative analysis tools will be used to identify themes in the transcripts, and to quantify key issues and what people thought about them. This will form the main part of the Discussion chapter in the thesis. Sometimes there may be a particularly salient point made by a participant, and I will get in touch to ask if you consent to this being added as a direct quote. No quotes or attribution of content to you individually will be made without your explicit consent.

The audio and transcripts will be kept in a secure location on AUT premises, and only the researcher will have access. These files will be kept in accordance with AUT Regulation (six years). You can request a copy of your data at any time.

What are the discomforts and risks?

The research is not on a particularly ‘controversial’ topic, and you are being invited to simply offer your thoughts around how decentralisation can work in practice – particularly within New Zealand. For participants currently employed in the New Zealand public sector, especially the health sector, there will be no questions around endorsing or supporting decentralisation, and participants will not be put in a position which may lead to conflict with current policy or strategic direction. The interview is much more of a ‘thought experiment’.

One other potential discomfort is around identification of participants. The value of this research stage is in summarising the views of relevant experts to build a picture of how decentralisation might work – specifically in New Zealand. The academic value of this is greatly enhanced when participants can be named, and it is clear which experts have contributed to the results. However, it remains your right to not be identified in this research if you so choose.

If you do not wish to be identified in the research then please consider the risk that any reference to your role or position (which would only be done with your agreement) may identify you indirectly – for example, if you are well known in a very specialist field. If you have any doubts, then please contact the Researcher to discuss.

E. Interview Participant Information Sheet

The Consent form will ask you to clarify whether or not you are happy to be identified.

How will these discomforts and risks be alleviated?

If you experience, or anticipate, any discomfort at all – arising from this research – then please contact myself or my supervisor as early as possible (details below).

What are the benefits?

In this research I want to highlight an important issue that is not well researched. Currently, there are initiatives to give people better access to their health data (for example, via patient portals). However, these tools do not give real *control* to people - there is a lot of research indicating that, if people have full control of their data, they feel more empowered and engaged with staying well. There are no current initiatives or research projects focusing on giving individuals this level of control.

The research therefore fills an important gap, as we move further and further into using technology to improve healthcare - all fuelled by the capture of data. The interview stage specifically offers a critical insight into how we can move towards decentralisation, something which has no identifiable research base yet, particularly in New Zealand.

The benefit to myself personally is that my research will lead to graduating as a PhD, and the findings may be presented in journal articles or at conferences.

How will my privacy be protected?

As already noted, participants are free to request they are not identified in the research. Furthermore, no direct quotes or attribution will be used without specifically seeking your consent to do so. The Consent form will ask you to clarify whether or not you are happy to be identified.

What are the costs of participating in this research?

The cost for participants is one hour of time to complete the interview. Participants are invited to review transcripts before they are utilised in the analysis, and this would require approximately one additional hour of time. The objective of the researcher is to make participation as quick, easy and convenient as possible so that your expert viewpoint can be included in the research.

What opportunity do I have to consider this invitation?

Purely due to scheduling issues, with overall completion of the thesis and Ethics Committee meetings, I would ask that you please respond to this invitation within one month of receipt.

Will I receive feedback on the results of this research?

Yes. You will receive a full transcript of the interview, and be invited to review it if you wish. Once the thesis has been completed, I will send you a summarised copy of the findings as well as a link to the final thesis for your reference.

What do I do if I have concerns about this research?

Any concerns regarding the nature of this project should be notified in the first instance to the Project Supervisor - Professor Marilyn Waring, marilyn.waring@aut.ac.nz, (+649) 921 9999 ext 9661.

Concerns regarding the conduct of the research should be notified to the Executive Secretary of AUTC, ethics@aut.ac.nz, (+649) 921 9999 ext 6038.

Whom do I contact for further information about this research?

Please keep this Information Sheet and a copy of the Consent Form for your future reference. You are also able to contact the research team as follows:

Researcher Contact Details:

Alex Poor, alex@robot5x.com, (+64 27) 297 4591.

Project Supervisor Contact Details:

Professor Marilyn Waring, marilyn.waring@aut.ac.nz, (+649) 921 9999 ext 9661.

Approved by the Auckland University of Technology Ethics Committee on **17 May 2021**, AUTC Reference number **21/146**.

F. Initial coding from Research Phase three

Presented below is a full list of all codes derived from the initial coding process of Research Phase three, as discussed in section 7.6.2. An excerpt for each is provided, however please note that these have been edited to protect the privacy of interview participants.

Table F.1.: Thematic Analysis: Initial coding

Code	Example
Analysing data from te ao Māori	<i>[Māori] might want to do their own analysis, or augment it with some other data, and they can't do that if they can't get the data.</i>
Decentralisation won't work	<i>I am not convinced that we would ever end up with fully decentralised everything.</i>
Decentralisation increases trust	<i>The idea that it's decentralised has a benefit in terms of helping us address those people that don't fundamentally trust their data being held centrally.</i>
Access to data improves health	<i>...or using that in a way where you go and investigate things about your condition and the data is there to support the ease of access to the information.</i>
Verified identities are important in sharing data	<i>When you're a verified, trusted identity, there's all this information that can be made available to you without any concerns about breaching privacy because it's your data.</i>
Conflating access to data with sovereignty	
Decentralisation reduces risk and liability	<i>[Decentralisation] takes the risk away from the central agencies of having responsibility for governing and protecting that data.</i>
Unreliability of source systems	<i>Some systems, I don't think we would want to trust as being accessible when we need them.</i>

Continued on next page ...

F. Initial coding from Research Phase three

Centralisation demands duplication of data	<i>The Ministry of Health shifts terabytes of data every week, and it's just a replica of what DHBs have already got so why do we need another copy of it?</i>
Decentralisation improves integrity	<i>The other benefit of decentralised data, which is that there's only one copy of it. I'm not trying to maintain the integrity of my copy of it.</i>
Ambiguity around Māori Data sovereignty	<i>One of the challenges of Māori data sovereignty is that people think it's synonymous with onshoring of data, and it's not. For some people, it absolutely is, but in the main, that's not what it's about at all. It's about ownership and control.</i>
Government need to access data	<i>In order for the system to work, the people who are trying to do population and public health need to have visibility of some of the data that might exist.</i>
Sharing data effectively requires control	<i>There was a push to open the IDI up because Māori wanted to use it for their own social policy. But when they realised that if you open it up it means Ngai Tahu might see data about Ngati Porou, they said "No, we're not happy about that."</i>
Decentralisation improved quality	<i>It's then your asset, up to you to maintain it. It shifts the balance of power to them.</i>
Domain legislation as an obstacle to decentralisation	<i>The Medicines Act would have some aspects that would probably prevent you from being fully decentralised</i>
Generational attitudes around social licence	<i>So, you could argue that social licence is not ubiquitous. It comes down to generations - Gen X, Gen Y or Baby Boomer.</i>
Need to see how your data is being used	<i>You should know all activity against your data. As a minimum - exactly who accessed your data.</i>
Clinical risk when access is controlled	<i>If you want your doctor to be able to make the best decision that they can about your health situation, then they probably need to have as complete a picture as possible.</i>
Decentralisation creates more opportunity for data	<i>I think you will see more pragmatic use of DNA and as you get into that as a consumer in healthcare, you want to own that data and control it yourself.</i>
Technical obstacles to decentralisation	<i>At a technical level, I would imagine that there would be quite a bit of work... we're talking about quite significant technical debt.</i>

Continued on next page ...

F. Initial coding from Research Phase three

Conflating decentralisation with blockchain	<i>Bitcoin is not a very efficient transactional processing environment, and so what would this concept of decentralisation mean from a transactional point of view</i>
People should have control over their data	<i>The principle of people having more control over their own data and being able to make decisions about how that data will be used is, in my mind, preferable.</i>
Control over data gives people agency	<i>The advantage is that it gives people agency</i>
Centralisation preferred because of easier access to data	<i>In the European context early last year, all of the public health entities wanted that health data to be centralised so that they can make the best, most efficient decisions possible with full oversight over the entire system</i>
Decentralisation introduces risk of not sharing data	<i>The Aotearoa New Zealand approach, which is the entirely decentralised approach, where all of the data remains on the devices, and you only provide that data to the Ministry of Health in the event that you get COVID-19 and you need to voluntarily upload that data. If you choose not to provide that data, then that significantly limits the Ministry of Health's ability to respond.</i>
Additional uses of centralised data	<i>In Australia, we have seen that police have requested access to that data for police purposes, not health purposes. In some states, the health authorities have declined those requests, and some states the health authorities have said, "Well, you've got a warrant so I guess we have to give you this data."</i>
Legislation currently protective of secondary uses of centralised data	
Social licence is not homogeneous	<i>You have to be genuinely open to the notion that maybe you can't do the thing you wanted to do. And that's something that decision makers don't always honestly engage with because they view the consultation part as maybe being more of a box-ticking exercise. They just say, "We held a hui, everybody heard about it, we can go ahead".</i>

Continued on next page ...

F. Initial coding from Research Phase three

Decentralisation depends on trust	<i>In a decentralised architecture, if I said to you, "Here's all my health data, I still own it, you can't give it to anyone else," if you then decide to do that and sell it on the dark web, there's probably very little that I could do to stop you. So there's a lot of trust that you're not going to do the wrong thing with that data, and that comes down to the fact that data is non-rivalrous.</i>
Decentralised identity management is very difficult	<i>In a decentralised architecture, it is really really difficult, and you have to just rely on encryption methods, like how I can prove who I am by using my PGP key. It's quite hard to really do that reliably.</i>
Data sovereignty is only a technical problem	<i>The challenges that they run into probably aren't so much legal, as they are technical and practical.</i>
Data sovereignty is geopolitical	<i>It is assumed that if the data is in a jurisdiction that you have control over, therefore you have control over that data.</i>
People will struggle with decentralisation	<i>If we were moving to a decentralised world tomorrow, so much education will be needed but people don't get it because they've been living in a centralised world for so long.</i>
Incumbents as an obstacle to decentralisation	<i>It's not in their interest for it to happen and so they will fight against it.</i>
Relationships are important in preparedness to share data	<i>When you can have a really good intelligent conversation, people are very willing to share their information if they feel that not only will they get a benefit from that, but also other people.</i>
We are missing out on data opportunities now because of consent	<i>There's hugely valuable health survey data that we cannot link and use because they never got the right consent in the first place. Because nobody thought about it, or we get scared.</i>
Right to be forgotten	<i>The very next day, the woman came in and said she wanted to withdraw everything she told them. The conversation did happen, right? That information is out there. What was she meant to do with that?</i>
'Positive' use of government data may not be welcomed	<i>Even though there is a benefit to citizens from using integrated data at an individual level, that doesn't make it right and it doesn't mean they're going to like the idea.</i>
Benefits of using data outweigh concerns	<i>The benefits outweigh the concerns around using people's information.</i>

Continued on next page ...

F. Initial coding from Research Phase three

Communities hold more powerful data than government	<i>There is hugely valuable information that sits in communities that will never ever sit in the administrative data that government collects and holds.</i>
Lack of capability around data in government	<i>When somebody snaps their fingers and says "set that up" and they have to set it up in 24 hours, the last thing they think about is how they might do the data piece. And in fact, even when they had more than 24 hours, the last thing they thought about was how they would do the data.</i>
Government have to demonstrate purpose in data collection	<i>They just did this blanket - we're just going to keep people's information without thinking about and really demonstrating why they needed the information they were collecting. And they couldn't make the argument.</i>
Data is about relationships	<i>The biggest thing in building trust is relationships and intention.</i>
Complexity of decentralisation is manageable	<i>We're doing all of this [permissions management] anyway. We're just doing it in a thousand different apps, passwords.</i>
Data can offer benefits which override privacy risks	<i>I usually argue for public good overriding the very small risk to personal privacy.</i>
Nobody has an absolute right to any data	<i>No one has an absolute right to any data. It's co-produced... so there's a bundle of rights to negotiate.</i>
Importance of standards	<i>[Federating data] also requires a very strong notion of data standards and metadata standards.</i>
Potential for abuse	<i>You will get some kind of bullying or other behaviour. There's potential for others to abuse peoples rights.</i>
Tension between group rights and individual rights	<i>I think the big cultural one in New Zealand is group rights versus individual rights. In particular, things like DNA.</i>
Problem of vendor lock-in	<i>They make it hard to get access to data. A lot of them are doing that. There's a lot of lock-in plays like that.</i>
It's not a technology problem	<i>It's going to take a lot more than just technology to sort it out.</i>
Growing mistrust between people and 'the government'	<i>That gets back to what our relationship with the government and the state is. It used to be very close in the early 20th century. People thought they were the government. Whereas now, I think it's further apart and more distant. There is less trust.</i>

Continued on next page ...

F. Initial coding from Research Phase three

Ownership of data is objectively different to other things	<i>Whereas with personal information, if I lend you my information, I still have it. I might let you go do what you want with it and return it to me, but I haven't given it away and I don't have any less of a claim to it than I did before I gave it to you.</i>
Is decentralisation easier in a public health system?	<i>How would that sit in a private health system where your insurance company might require you to give it all to them because otherwise they won't insure you.</i>
Can't guarantee integrity of patient-generated health data	<i>Another [issue is] ... the reliability and accuracy of apps that collect information, which leads people to either believe that they're healthy when they aren't, or causing them to push their body... In a decentralised system, your ability to control that might be compromised.</i>
How do we protect access to decentralised apps?	<i>If you look over at bitcoin, cryptocurrency, blockchain systems, is what happens if you lose your key? You've lost your money. How do you access this information in an emergency?</i>
We need to rethink consent	<i>Consent has to be contextual. People have to be able to have the ability at any point to rescind that consent. Because if they don't have that mechanism, then why would they give consent in the first place?</i>
We don't need technology solutions for trust or data sovereignty	<i>So when you look at organisations that spend the time and do those things right, I don't think they need different technology solutions to support them to do the right thing because I think they're already doing it.</i>
Lack of a clear strategy about the steps we're going to take to achieve outcomes	<i>Why does the Ministry of Health not have a data strategy that says, "...in an ideal world we would have this information because it would support us to do these other things we currently can't do, and the benefit to the people of Aotearoa is x, y and z. Over the next 10 years, we're going to take some deliberate steps to move ourselves towards collecting that information."</i>
Government has not made its data accessible to people	<i>"How do we get data out of the IDI and into the hands of communities?" This is sovereignty. This is not about who owns the data, it's about whether or not I can use it to make decisions to inform my decisions.</i>

Continued on next page ...

F. Initial coding from Research Phase three

More important that data is used for good than who it belongs to	<i>Part of it is about where data belongs, but it's less important to me that it belongs to me, it's more important to me that the people who it belongs to are receiving benefit from its use.</i>
Government does not have an incentive to be agile and demonstrate value with data	<i>The approach in the corporate world is to go and try something small and test it and add value, rather than the government approach to just spend millions of dollars on big things without stepping back and thinking about the outcome they're seeking.</i>
Personalised medicine requires much broader data collection than present	<i>The data that they use about us as patients is going to have to grow and therefore be better than it is today, and widely available. It has to be more clearly owned and controlled than it is today.</i>
Philosophical data sovereignty is more important than geopolitical data sovereignty	<i>For me as a patient, I'm not worried about where it's stored, what I'm worried about is that it's owned by me, that I know where it is, what it is and who has access to that.</i>
Centralised models cannot handle data requirements for next generation	<i>I just don't think centralised is going to be able to handle what's going to be required.</i>
Decentralisation is aligned with Māori data sovereignty	<i>You're looking at decentralising the whole system, and that is aligned with Māori data sovereignty.</i>
Difficult to imagine decentralisation because centralisation is so dominant and we don't understand the possibilities	<i>I think we don't know the full extent of what the opportunities are yet, because we don't know that this world exists.</i>
NGOs/providers working with hard to reach communities have problems accessing data currently	<i>Accessing basic aggregate public health data is really difficult... As a proponent of open data and being able to share data safely, we shouldn't even be in this position.</i>
Iwi should be enabled to make most of opportunities with data	<i>Māori people are going to trust iwi much more than they will trust the government when collecting data.</i>
Decentralisation could develop separately in the provider space	<i>"You guys collect this stuff, and we'll now focus our community efforts on this stuff because this stuff is important to us." If we build the decentralised system properly, the government would ideally support this too.</i>

Continued on next page ...

F. Initial coding from Research Phase three

Consent is working poorly because of general lack of understanding/capability on it	<i>With consent, it's really difficult. When it comes to data, when you're trying to get consent from someone who doesn't understand it, there is an education process that has to come along with it.</i>
Very low capability around data in people generally	<i>They do care about data and information, it's just that for the majority of everyday people, data isn't a language that they talk in. Actually, when you talk to communities and show them what it is, they're very interested and do care.</i>
Radical changes in public sector are difficult because of funding and change management issues	<i>They're generally not making the right decisions based on what's good for everyone. The lens becomes, "What can we do with the money that we've got?"</i>
Decentralisation introduces real clarity and assurance around data relationships	<i>There potentially are some advantages for agencies or others who are users of data, in that if someone does give you access to that data, then you are very clear what kind of scope of use and what is appropriate to be done with that data.</i>
Having enough representative data for evaluation in decentralisation may be tough	<i>If you need a representative sample, can you get enough people in all of the different groups that you would need to get?</i>
Legal framework not currently supportive of collective rights around data	<i>A third-party collective that has the interests of their members is not well catered for by the current system.</i>
Public need better understanding of current data flows	<i>I think if we did it well, we would build social licence, because the types of things you hear from the DPUP work, as you say, are, it's not that we're overly bothered that the system needs to use information, but we do expect to be told about what it's being used for and why it's being used that way.</i>
Benefits of decentralisation could be polarised by engagement and capability	<i>Maybe from an equity of outcomes perspective, we would end up with people who are engaged with it more getting more value from it, but people who didn't engage with it quite so much seeing problems.</i>
Lack of regulation/oversight has enabled vendor power around data	<i>If you hold the data, you've sort of been allowed to come to the point of view that you can control the data, even though I don't think that's an explicit setting. I just think that because nothing's been said, people can just form their own view basically.</i>
Security and privacy build trust	<i>Part of the trust is that it's secure and that it's private.</i>

Continued on next page ...

F. Initial coding from Research Phase three

Access to data is a product of trust	<i>There's the notion of trust that the more trust you give, the more trust you get. But people don't want to give up control.</i>
Engagement is a product of control	<i>How do you get people to engage in that? You tell them they have control.</i>
Centralisation is convenient	<i>If they could go from one place and control it all from there, it'd be great. That's not the landscape at the moment.</i>
Geopolitical data sovereignty is a distraction from the centre	<i>I think the health sector engineers the conversation to focus on where data is hosted.</i>
Digital equity	<i>There's a big challenge around not creating on top of health inequity, digital inequity.</i>

End