

—
**THE NHS FEDERATED
DATA PLATFORM AND
PALANTIR**

—
7 KEY RISKS

—
June 2023

A report by



CONTENTS

04 Executive Summary

08 Introduction

Background: the NHS Federated Data Platform

A brief history of Palantir, Foundry and the NHS

12 Risk 1: a flawed procurement process

14 Risk 2: a vague and growing project scope

16 Risk 3: secrecy and a failure to design for patient consent

18 Risk 4: over-centralisation

20 Risk 5: monopoly lock-in

22 Risk 6: Palantir's poor reputation

21 Risk 7: failed pilots

24 Alternatives

26 Recommendations for policymakers and legislators

27 Conclusion

28 Endnotes

31 Acknowledgements

EXECUTIVE SUMMARY

England's health and social care systems are in crisis. The strains of the pandemic and a huge backlog of operations are compounded by chronic underinvestment, staff shortages, and industrial unrest.

Whilst tech will never replace the need for properly paid staff and sustainable budgets, improved use of data and technology could undoubtedly help the NHS meet future needs, improve patient care and increase efficiency. Past efforts to harness the power of NHS data have repeatedly failed, principally because public concerns about private sector involvement and privacy have not been addressed.

Right now, the NHS has over 3,000 vacant tech roles; its vital data science capacity is being actively slashed.¹

Instead of investing in its own tech and data science capacity, NHS England is currently preparing to pay Palantir £480 million to build a "Federated Data Platform". Palantir, a US tech firm which had no track record in healthcare, is in pole position to win this contract – largely as a result of pandemic opportunism and a concerted lobbying effort. If Palantir wins the contract, its proprietary software, Foundry, will be permanently installed in the NHS, along with numerous Palantir consultants the NHS will have to retain to run it.

This report identifies seven key risks with the NHS Palantir deal:

A flawed procurement process – the frontrunner, Palantir, is in pole position largely thanks to (a) pandemic opportunism, which unfairly secured incumbent advantage when procurement rules were suspended, and (b) a problematic lobbying operation.

A poorly defined and shifting project scope – this bears many hallmarks of past government procurement catastrophes, where a flawed understanding of how, and if, a programme will be used wastes billions in public money.

Secrecy and a failure to design for patient consent – As with earlier projects, excessive secrecy risks fatally undermining trust in the FDP. High levels of mistrust in health data plans are dangerous, because the only practical way patients can dissent is to trigger a legacy 'opt out'. The National Data Opt – Out is a blunt instrument – throwing out positive uses of data along with more controversial ones. A lack of trust has left opt-out levels extremely high, at over 3 million.

Recent polling from YouGov has found almost half of adults in England who have not yet opted out, 48%, are likely to do so should the Federated Data Platform be introduced and run by a private company.² If anything like that number actually choose to exercise their National Data Opt-Out, it will have a catastrophic impact on the quality of NHS data, an extremely precious British health resource.

48%

of adults in England who have not yet opted out are likely to do so should the Federated Data Platform be introduced and run by a private company.

30% very likely **18%** fairly likely

Source: YouGov

Over-centralisation – exacerbates public trust issues, creates needless coordination challenges and security risks, and stifles local and regional innovation.

Monopoly lock-in – would entrench NHS dependency on Palantir, obstruct integration with other NHS systems, and obstruct innovation on other platforms.

Palantir’s poor reputation – its long association with US security and surveillance operations, and a Trump-supporting, NHS-bashing owner, risks further eroding trust amongst the public and NHS professionals.

Pilot failures being ignored—there is a lack of transparency around the data or documentation used for claims that uses of Foundry have so far been a success. NHS professionals are telling us off-the-record that significant problems are being glossed over and many pilots have been suspended for failure to work.

Whilst there is a lot of momentum behind NHS England’s deal with Palantir, it is not too late to change course. Many existing initiatives offer alternative visions of how the NHS could harness data. A consortium of UK-based companies developed their own alternative, Palantir-free bid for the FDP contracts, but never had a real chance to compete. NHS data scientists, leaders, and experts have put forward credible and cost-effective alternatives to help the NHS use patient data safely for care, planning, and research. Some NHS regions, like London, already share some data to plan care. The openSafely system from Oxford, developed in the pandemic, already enables safe research.

The best alternatives to the FDP keep power and access at a regional level, so decision-making stays close to the communities who rely on the NHS. They use open-source solutions, which keep the value of innovation with the service, and they honour patient trust.

If the government and NHS England plough ahead with the FDP process without addressing the risks we have identified, the project risks either unravelling entirely or becoming an expensive failure which acts as a drag on NHS performance and finances for years to come. We hope this report helps encourage a reset.

DAUK and Foxglove's recommendations:

Parliament should investigate the FDP procurement: its design, its value for money, the risks the system poses to patient trust, and whether an adequate procedure has been followed to ensure all bidders are treated equally and fairly, and are fit and proper partners for the NHS.

The FDP tender should be paused until these questions are answered. If no satisfactory answers are available, the tender should be withdrawn.

The government should consult on and design alternatives for managing patient data that better use existing NHS capacity and taxpayer funds. If external support is required, any system should avoid locking in the NHS into one monopoly provider.

The government should reform how patients can consent to (or opt out of) data-sharing, giving patients a more granular say in how their health record is used.

INTRODUCTION

—
The fact that so many committed health professionals have concerns about the project, but do not feel able to share these concerns publicly, is itself concerning.

This report lays out what DAUK and Foxglove have been able to uncover about the plans for an NHS Federated Data Platform, how Palantir came to be the clear frontrunner for the contract, and lays out 7 key risks with the deal.

It draws together information from a range of different sources, including information published by the government and the NHS, information we have gained challenging Palantir’s involvement, including in the courts, and information which has been shared with us in confidence by a range of NHS professionals.

A significant number of NHS insiders have assisted the research for this report, but all of them only felt able to speak “*off the record*”. The fact that so many committed health professionals have concerns about the project, but do not feel able to share these concerns publicly, is itself concerning. The National Audit Office highlights a “*lack of honesty*” as a key reason why major government projects go off track and fail to reset in a timely manner. That NHS officials and data professionals feel more able to be candid to DAUK and Foxglove than to those driving this project should ring alarm bells.³

There is still time to change course. Many NHS professionals we spoke to expressed frustration that superior alternatives are available, but were not considered. We hope this report will encourage NHS leaders to reconsider their approach, and give parliamentarians reason to insist on a pause whilst they investigate this process and the alternatives fully.

Background: the NHS Federated Data Platform

The NHS holds a huge amount of health data. In fact, it's one of the most complete sources of longitudinal health data in the world. The data is so good thanks to the fundamental values of the NHS as a health service that cares for everyone, which means the data covers all of us, across the UK, regardless of wealth or background.

Responsible use of this data can be of huge public benefit. It could enable significant improvements in clinical practice, drug development, and healthcare service planning. Patient data helped the response to Covid, and, with proper safeguards to protect the public interest, more could be done in the future. One estimate of the value of patient data to the NHS puts it at a staggering £10 billion per year.⁴

There have been several previous attempts by the NHS to develop more integrated data systems, all of which have collapsed because of a lack of public trust. Care.data was a programme, launched in 2013, to transfer all patient data from GP surgeries to a centralised database. This aimed to support resource planning and research to improve care. Plans included linking the data to other NHS datasets and offering potential to external parties, including academic researchers and private companies.⁵

A lack of consultation, a poorly designed opt-out process, and the involvement of private companies caused widespread concern. Controversy led to a series of delays, with the project formally “*paused*” twice, in both 2014 and 2015, in attempts to placate public concern without addressing the fundamental issues with the plan. Mistrust of the scheme remained high and huge numbers of patients opted out of their data being shared. The level of opt-out, some 1.2 million patients, degraded the quality of the data and was

a major reason that in 2016 the NHS was forced to abandon the programme.⁶

In 2021, this failure repeated itself. Another system, ‘General Practice Data for Planning and Research,’ (GPDPR), sought to extract patient data from GP records into a central government database. Again, officials failed to satisfy the public that there would be safeguards for patient data against collateral misuse. Again, the plans triggered a mass opt-out wave – over a million in a month alone. Again, the project was shelved.⁷

The “*Federated Data Platform*” is just the latest in a line of ill-fated NHS national data integration projects. Currently under an active procurement process, with the first contract reportedly worth £480 million, the FDP would be the largest single point of access to NHS patient data ever created. This procurement is said to be the first of four, making the system’s ultimate cost unknown.⁸ The front-runner to run this system has long been known to be the US tech company Palantir, and their “*Foundry*” product.

Experts have described the FDP as “*the operating system for the NHS,*”⁹ or “*the veins through which patient data will flow.*”¹⁰ The FDP prospectus is worryingly vague, but in short, the FDP aims to create a single access point to patient data for hospitals, GPs, and social care – all to be accessible to central government on a single software platform. The data will be pseudonymised, but would nonetheless be extremely sensitive patient data.

In principle, the joining up of patient data to benefit the NHS is a sound and long-overdue idea. But for the FDP plan to succeed where previous schemes have foundered, it would need to acknowledge and solve the major problems which ultimately derailed previous projects.

Background: Palantir and its Foundry software

Palantir was founded in 2003 by a Silicon Valley billionaire called Peter Thiel, who is the current chairman and thought to be its largest shareholder. It is headquartered in Colorado.

Originally funded by the CIA, Palantir's core business is providing big data and surveillance support to military, security, intelligence and police agencies. Its clients include the US military, the US National Security Agency (NSA), the US Immigration and Customs Enforcement (ICE) agency, and various US police forces. In 2020 it won a contract from the UK Home Office as part of the UK's post-Brexit border enforcement tool.¹¹

As you'd expect from a company specialising in security and surveillance, much of Palantir's work is secret. However, what we do know reveals a pattern of involvement in controversial programmes criticised for abusing poor people, migrants, and other vulnerable groups:

Family separations and migrant deportations: Palantir was a key enabler of Donald Trump's extreme anti-immigration policies. ICE used Palantir to run workplace deportation raids and seize family members of unaccompanied migrant children.

Mass surveillance by the NSA and GCHQ: Palantir helped the US and UK's digital spy agencies (NSA and GCHQ) manage mass surveillance programmes like XKEYSCORE. This was one of the systems exposed by whistleblower Edward Snowden for tracking millions of innocent people's movements online.

Racist policing in the US: Many US police forces use Palantir for "predictive policing" – widely criticised for unfair targeting of poor and black communities. In LA, the police used Palantir to build a tool, "LASER", that claimed it would extract suspected offenders from the community "like a tumor".

Secret Operations in Iraq and Afghanistan: Palantir provided intelligence software to power the Iraq and Afghanistan wars.¹² The exact nature of their involvement remains secret, appearing to involve several overlapping but distinct uses of the software.

"Foundry" is one of Palantir's key products. It is essentially a data platform, designed to bring together disparate sources and types of data into a single, accessible platform. Its marketing materials claim that it is more than that, "*an operating system that coordinates the interplay of data, models, and decisions in an enterprise*". It was originally a corporate-facing fork of Palantir's "Gotham" system, which is more associated with security and policing uses including U.S. Immigration and Customs Enforcement's (ICE).

Customers for Foundry have historically mainly been private sector clients in sectors such as aerospace (such as Airbus), automotive (such as Ferrari) and oil and gas (such as BP, ExxonMobil and Trafigura). Palantir now seeks to market itself to the health sector, but the case studies and partners lists on this section of its website are strikingly thin. It claims to be "*relied upon by providers, healthcare organisations, and healthcare systems such as the UK's National Health Service (NHS)*" but the NHS is the only example offered.¹³

Palantir secured its first contract from the NHS to provide Foundry in 2020, at the height of the pandemic, when ordinary procurement policies were suspended. Foundry was selected to provide an emergency Covid-19 datastore, with an initial contract at a nominal £1. Its use has been extended through a series of short-term contracts, the most recent of which was worth £11.5 million.¹⁴

As you'd expect from a company specialising in security and surveillance, much of Palantir's work is secret.

A FLAWED PROCUREMENT PROCESS, MARRED BY LOBBYING AND INSIDER DEALING

Government procurements must be transparent and scrupulously fair. This FDP procurement process fails at this first hurdle – it is the culmination of years of influence peddling by Palantir, exacerbated by the waiving of ordinary procurement procedures during the pandemic.

Prior to gaining a toe-hold in the NHS, at a time when procurement rules had been suspended, Palantir had no track record in health. Palantir is mainly an intelligence and defence contractor whose systems support surveillance and policing. Private sector use of Foundry has been concentrated in very different sectors such as aerospace and energy.

However, by 2019 Palantir had identified healthcare as a market it wished to move into and the NHS as a target. So it started lobbying government officials. In early 2019, Liam Fox and officials in the Trade Department held closed-door meetings with Palantir bosses at Davos. Briefing notes say these meetings were set up to discuss the ‘untapped’ commercial potential of NHS data¹⁵ – and they kicked off a charm offensive by Palantir that

ultimately helped get the firm’s toe in the NHS’s door.

In mid-2019, Palantir’s Louis Mosley wooed the then-chair of NHS England’s board, Lord David Prior, over watermelon cocktails at a closed-door dinner.¹⁶ Over the coming months, Palantir executives heavily lobbied Prior, Matthew Gould, and other senior NHS officials. Prior flew to California in January 2020 for private meetings with Palantir.¹⁷

For Palantir’s NHS aspirations, the pandemic therefore came at an extremely opportune moment. The Covid emergency meant the usual procurement rules were suspended, and NHS leaders under extreme pressure to adapt the health service to an unprecedented challenge.¹⁸ Palantir landed its crucial first data management contract with the NHS by offering Foundry to run a vast new ‘datastore’, for just £1.¹⁹ There was no competitive tender for this contract; the deal was simply handed to Palantir.

At the contract’s first renewal, Palantir increased the price of its services to £1m. On its second reboot, the price went up to

—

Briefing notes say these meetings were set up to discuss the ‘untapped’ commercial potential of NHS data.

£23m. It has since won a further £11.5m extension. In all, Palantir used a no-bid deal at the start of the pandemic to win at least £34m of public funds.

There have been other high-profile examples of the costly results of closed-door pandemic emergency contracts. But whilst the sums wasted in the procurement of PPE have so far received the most publicity,²⁰ the £1 deal which established Palantir’s beachhead in the NHS – without meaningful scrutiny, competitive tender, or debate – could end up having more profound and costly long-term implications.

This is because for Palantir the most valuable outcome of these Covid contracts was not the £34m, but the opportunity to install itself, without any competitive tender process, in the heart of the NHS.²¹ This gives Palantir an obvious incumbent advantage in bidding for the contract to provide the NHS with a permanent system – the FDP – and a contract which thus far has risen in value to £480 million.

Faced with public opposition to the firm’s role in the NHS, Palantir has continued lobbying. In parallel it also ran a second strategy, which UK chief Louis Mosley indiscreetly described as, “*buying our way in*” by “*hoovering up*” small businesses serving the NHS to “*take a lot of ground and take down a lot of political resistance.*”²²

Palantir’s strategy has worked. NHSx’s original public statements about the

Covid datastore promised that it was an ‘emergency system’, and that it would be unwound at the end of the pandemic to make room for a more permanent solution, based on patient trust.²³ Those promises have been quietly withdrawn. One of the authors of the original blog which made these pledges has since taken a paid position at Palantir.²⁴

The FDP is a very consequential contract for the future of the NHS: it is envisaged as serving as core NHS data infrastructure for years to come.

Palantir, a middling data analytics firm with no track record in health, is in line to become the ‘*operating system for the entire NHS*’.²⁵ The contract is designed in a way that makes it almost impossible for anyone but Palantir to win. The tender contains phrases that are taken from Palantir’s internal systems.²⁶

Home-grown UK health tech firms would have liked a fair chance to bid – and one UK consortium said it can do the job for a fraction of the cost.²⁷ But they are struggled to compete with Palantir’s colossal incumbent advantage, won by insider influence, pandemic opportunism, and lobbying.

A VAGUE AND GROWING PROJECT SCOPE

The National Audit Office has previously highlighted that problems with scope are a key reason for major public sector projects going off track and over-budget.

In its 2020 analysis, *Lessons Learnt From Major Programmes*, it observes that “a flawed understanding of how a programme will be used in practice can lead to gaps in its scope, requiring costly changes later on, and also to bodies making programme decisions with potentially far-reaching unintended consequences.”²⁸

Sadly these “*Lessons Learnt*” are in danger of being ignored with the Federated Data Platform.

The FDP prospectus has proven difficult for even experts to follow. It seeks to create a single point of access to hospital, GP, and social care records. There are five stated uses in the initial tender,²⁹ but it is said that these will expand.³⁰ In other words, we don’t know exactly what the government seeks to do with this vast new database.

The documents also state the customer base for the FDP is “wider than just NHS England,”³¹ without clearly explaining who else will have access. Integrated care systems will have access, but so too, it is suggested, will providers commissioned under a contract to provide NHS services, which potentially includes a raft of public and private organisations.

A lack of precision in project scope increases the likelihood of the NHS getting

poor value for money. It increases the risk of being left with a system which doesn’t fulfil the NHS’s core needs, and of paying for features it doesn’t need at all. Changes to the scope later in the project are liable to incur increased costs.

We haven’t been told why the value of the FDP contract keeps rising, from £240m to £360m to £480m – with the procurement said to be the first of four, but this would be typical of a project with a vague and expanding scope.³²

These concerns are sharpened by recent statements that each NHS trust will be given their own separate instance of the platform - meaning the cost could run into the billions.³³

This looseness around what the project will do also exacerbates the risk, explored in more detail below, that the project can be fatally undermined by public mistrust. This is because the vague scope leaves potential for the FDP to enable the re-use of patients’ sensitive data in ways that exceed public understanding or consent.

For example, the FDP envisages health data being processed “outside the UK”³⁴ as well as a “marketplace” which the government hopes will “enable the development of new and novel data driven solutions – e.g. apps – that can be made available to end users.” Few patients now think, when they go to the GP, that their health record may be used to develop an app.³⁵

In other words, we don't know exactly what the government seeks to do with this vast new database.

SECRECY AND A FAILURE TO DESIGN FOR PATIENT CONSENT

One of the greatest risks of this system is that a failure to explain use cases and a failure to win over the public may fatally undermine public confidence in the system. This poses risks greater than the FDP – because it risks triggering yet a further wave of opt-outs that, because of the extremely blunt nature of the legacy opt-out system, undermine data quality in the NHS as a whole.

This has happened repeatedly in the past. Previous NHS data projects have foundered because of a failure to win patient trust. By the time care.data was abandoned in 2016, over 1.6 million NHS patients had opted out of the scheme – enough to seriously degrade the quality of the data. The National Audit Office found that there had been “insufficient focus on how to gain patient acceptance to share their data”.³⁶ In 2021, a further programme called General Practice Data for Planning and Research was put on hold after over a million patients in a month opted out of sharing data because they didn’t trust the government’s latest plans for their health record.³⁷

That means the current number of opt-outs, via the NHS’s “National Data Opt-Out” (NDOO) scheme, has grown to 3.3m.³⁸ This is a statistically significant

level of opt-out, negatively impacting the quality of the data and the accuracy of analysis which relies on it.

New research from YouGov commissioned for this report shows that **nearly half of English adults who have not yet opted out, 48%, are likely to do so if the Federated Data Platform is introduced and run by a private company.** Out of that 48% figure, the amount who say they are “very likely” to opt out, at nearly a third, or 30%, is almost double the number who say they are “fairly likely” to opt out, on 18%. That demonstrates the strength of feeling in the country rejecting a Federated Data Platform controlled by a company like Palantir.

All political parties agree that correct use of NHS data in future is essential to the continued existence of the health service. If anything like the millions of people reflected in YouGov’s study choose to opt out of sharing their health data, it will threaten the long-term usefulness of NHS data for everyone. **The 48% polling suggests a risk of over 20 million further opt outs – which would be a catastrophe. Even a fraction of this result is a major risk to the NHS’s data quality.**

If the government has done its own analysis looking at the likely number of opt-outs should the Federated Data Platform be introduced, with or without Palantir, they should publish it immediately. If no analysis has been done, one urgently should be.

There is particular public sensitivity to questions of privacy and consent concerning health data, because trust between doctor and patient is so essential to effective treatment. The system only works if patients believe what they tell doctors is private. Concern for the privacy of patients' health records has been repeatedly shown to run much higher than in other areas – because the data is sensitive and can cover areas like sexual abuse, or reproductive problems, or addiction.³⁹ For some patients, another expressed concern may be government overreach into health data. The former head of NHS Digital, Kingsley Manning, resigned over the Home Office repeatedly demanding migrants' data⁴⁰ – making the concern all the greater with a new health data store of unprecedented size and scope.⁴¹

This sensitivity is exacerbated by the extremely ham-fisted nature of the available opt-outs.⁴² The legacy systems for registering patient dissent – the Type 1 and National Data Opt Outs – are blunt instruments. The National Data Opt Out is the only way patients can register any objection to the sharing of their data by the FDP. But this tool only allows patients to opt IN to everything, including secondary uses they may not support⁴³ or opt OUT of everything, including secondary uses they would support.

An improved consent regime would give patients more granular control. A recent Imperial College report found that the entire NHS patient consent system badly needs reform, and a new consent system should give patients more information and control, including, for example, opt-out choices by data customer, not just by use.⁴⁴

To improve the data, and ensure the viability of the FDP project, trust will need to be won back and patients who have previously opted out persuaded to opt back in. That requires learning from past failures which prompted so many patients to opt out, and developing a system based on trust, consent, and transparency. Yet the FDP appears to be doubling down on past mistakes, and seeking to dodge rather than address fundamental issues with patient consent. It is on course to encounter similar levels of distrust to its predecessors, with similar implications for patients choosing to opt-out and dire consequences for the quality of the data. MedConfidential, a non-profit organisation which played a key role in the previous two waves of opt-outs, recently described the FDP tender process as a “shambles”.⁴⁵

Far from being likely to persuade patients who are currently opted out to opt back in to their data being shared, the FDP is currently on track to prompt a fresh wave of opt-outs. This poses such a huge threat to the quality of the data, and therefore the viability of the whole project, that it is hard to understand how this approach is rational or reasonable.

If the huge risks posed by a fresh wave of public opt-outs were not enough, the FDP's patient consent regime is also at risk of legal challenge for non-compliance with privacy laws. The FDP envisages radical changes to how patients' data is used and who can access it.

NHS leaders have described these changes as “*transformational*”, so they can hardly claim that they aren't significant. It is hard to see how making such changes to how a data subject's data is used, without explicit consent from the data subject for that change of use, is compatible with the requirements of the GDPR and UKDPA or the common law duty of confidentiality.

OVER-CENTRALISATION

The FDP also represents a major push to nationalise access to health data. While in some areas this nationalisation may be justified, it creates serious risks.

A national point of access is likely to be important for functions like vaccination monitoring, general planning, or pandemic response.

However, centralisation of the health data of over 55 million English NHS patients also carries significant risks, and should be done to the minimum extent necessary. One risk is collateral government use. Previous examples of undisclosed deals with other government departments (such as the Home Office and Immigration Enforcement Services) show that centralisation lends itself to government overreach.⁴⁶ This, in turn, creates risks that people who fear government access to their health records do not access care at all. This can have knock-on effects for public health, as for example people are not treated for communicable diseases like TB or immunised against Covid.

Another major risk associated with over-centralisation is security. The more data that is accessible from a single national point, the more catastrophic any potential data breaches would be. This interacts, in obvious ways, with the trust issues already explored.

Another risk is that an over-centralised system introduces unnecessary inflexibility. The present design of the FDP ignores the fact that many categories

of NHS decision-making, and most provision of healthcare, happen at the local or regional level. Local and regional innovation to respond to local and regional circumstances and needs will be stifled, or systems needlessly duplicated at extra cost to the taxpayer. The Imperial study mapping data flows observes, correctly, that there already exist extremely large pooled national health datasets, and that this procurement will simply generate an analytics infrastructure to sit atop them.⁴⁷ This seems duplicative, and raises questions about the project's value for money.

NHS managers who informed the drafting of this report note that, thanks to technological advances, these problems could be resolved by preserving more authority regionally, with regional approvals for use of the data and the analytics systems to be used, including members of the public. It would then be possible to link data nationally, but only if regional managers agreed on the value of the application, rather than delegating important decisions to an unknown, opaque central committee.

Keeping more power of data sharing at the regional level would have the benefit of building on existing practice. Some regional levels are already sharing data for health, planning, and social care. OneLondon, for example, already uses a shared data environment for some 9 million Londoners.⁴⁸ It recently awarded a data management contract to Better for just £3m for several years – that is, less

than 1% of the FDP's tender price to serve almost 20% of the same population.⁴⁹ What Foundry would add to this system, for manyfold the expense, is far from clear.

Finally, centralisation also raises more general questions about governance – such as who will decide about access for other future secondary uses, like research. Research is not listed in any of the five initial use cases for FDP, and government say it is not currently part of the initial plans. But research has been repeatedly described by senior officials as a core aim of any data centralisation, and is an obvious candidate for future use of the FDP architecture. In a Jan 2023 article talking up the potential of the FDP, NHS Ming Tang, Chief Data and Analytics Officer, NHS England, and Tim Ferris, National Director of Transformation, highlighted how one of the FDP's immediate predecessors, had “*informed research into treatments that saved lives*”.⁵⁰ Improving health research is undoubtedly a positive, but most NHS patients still wish to have a direct say over when, and if, their health data are used in research, and may have legitimately differing views about participating in for-profit research.

Instead of overriding consent and creating a single point of access to patient data, the government should be distributing resource and decision-making authority to the best regional examples of innovation across the NHS.

The present design of the FDP ignores the fact that many categories of NHS decision-making, and most provision of healthcare, happen at the local or regional level.

MONOPOLY LOCK-IN

The other great risk of a centralised national platform is that it will lock in the NHS to a single monopoly provider, which raises serious financial risk. Best practice for health data management would involve a) a blend of suppliers, so the NHS is not overly beholden to a single firm,⁵¹ and b) a system that the NHS can readily move on from. For the NHS to secure value for money, Foundry must be interoperable with other current NHS systems, and other future systems including potential alternative suppliers.

It should be easy, in other words, for the NHS to use Foundry alongside other systems such as OneLondon – mentioned above – and for the NHS to switch to an alternative to Foundry once any contract with Palantir has expired.

However, the data scientists and NHS managers we have spoken to report this is the opposite of the case with Palantir. By technical design Palantir has explicitly avoided creating a system that you can easily extract data from or that works well with other industry-standard data analytics systems. Palantir's system pushes people to its own proprietary systems; and switching costs will be very high.

For example, most NHS data scientists work in Jupyter, an industry standard data science tool. Palantir offers its own proprietary tool that does not work with this system. Other insiders report that another pilot of a central government use of Foundry, unrelated to health, took months to have the data

extracted. Palantir repeatedly promises interoperability, but NHS professionals with experience of Foundry have reported to us that they do not deliver.

This appears to be consistent with the experience of other public sector clients of Palantir. For example, US police departments have found that *“once you sign on with Palantir, it can be hard to sign off”*.⁵²

Incompatibility with other systems, opaque pricing, poor specification of needs from the commissioning force, and a conscious effort from Palantir to foster dependency on its consultants (for example claiming that only it could provide maintenance to *“proprietary”* systems which were in fact *“off-the-shelf”*) all seem designed to reduce its customers' ability to use Palantir products alongside alternatives or to consider competitors.

This is no accident. Rather, it is a deliberate design decision from a company seeking to entrench itself as a monopoly provider. Palantir CEO, Alex Karp, has explicitly said Palantir is seeking a monopoly position, explaining to shareholders in August 2022 that *“We are working towards a future where all large institutions in the United States and its allies abroad are running significant segments of their operations, if not their operations as a whole, on Palantir. Most other companies are targeting small segments of the market. We see and intend to capture the whole.”*⁵³

In the shorter-to-medium term, adopting Palantir would be likely to obstruct integration with other NHS systems and processes. It is also likely to hinder innovation on other platforms, including innovation which might be necessary at the local or regional level to make up for the over-centralisation and inflexibility of Palantir's system.

In the longer term, there's a risk of the NHS becoming locked in, with its dependency on Foundry and the atrophying of its capacity to use alternatives meaning it becomes ever-harder to hold a truly competitive tender once Palantir's contract has expired. The bargaining power of the NHS will get even weaker, and the cost of the services is likely to increase.

By technical design Palantir has explicitly avoided creating a system that you can easily extract data from or that works well with other industry-standard data analytics systems.

PALANTIR'S POOR REPUTATION

Previous NHS data projects have failed because of a lack of public trust. Palantir's poor public reputation, as a company associated with military, security, and policing, and with a Trump-backing, NHS-bashing founder chair, is likely to exacerbate these concerns. However, the level of reputational risk posed by Palantir's association with the NHS goes beyond the FDP.

During the pandemic the government ran public awareness campaigns encouraging migrants to get vaccinated – by promising not to share this health record with the Home Office for immigration enforcement purposes as they had previously done.⁵⁴ This reflected a recognition that amongst key communities, distrust was a barrier to accessing healthcare, and therefore a threat to public health.

Installing Palantir in the heart of NHS data infrastructure will exacerbate these trust problems. Unlike NHS Digital, the statutory safe haven for patient data which was recently abolished⁵⁵ and which previously managed national health data sets, Palantir have not been set up to be a check and balance on the government.⁵⁶ Palantir exists to do what its government customers ask of it – full stop.

Predictably, given this approach, other state uses of Palantir have recently been found unconstitutional by national European courts. Germany's Constitutional Court recently found the use of Palantir's surveillance software in predictive policing unconstitutional.⁵⁷

These concerns are amplified in the context of health data. While it may not be Palantir's

responsibility to ensure HMG uses health data ethically, it is the responsibility of government to do so. The sheer size of the FDP, paired with a supplier who has repeatedly facilitated unlawful government conduct, risks undermining confidence in the NHS.

Trust in Palantir is also not bolstered by their own executives' explicit celebration of privatisation, warmongering, profiteering, and their stated desire for monopoly lock-in. These statements begin at the very top – with Palantir's chair, Trump donor Peter Thiel. At a January address in Oxford he said that British love for the NHS was "*Stockholm Syndrome*" and added: "*In theory, you just rip the whole thing from the ground and start over*".⁵⁸

CEO Alex Karp's statements on pursuing monopolistic state contracts are canvassed above. Palantir's CTO Shyam Shankar (who has previously stated he hoped to see Palantir in "*every missile and every drone*")⁵⁹ recently wrote an article arguing that colossal profit margins are critical for defense contractors like Palantir – and that governments should accept this:

The problem with **defense contracting is not the popular narrative that contractors make too much money, it is actually that they make too little money... If innovators can provide capability for less money, why does the government care what the profit margin is?** Innovators will need outsized profits to motivate progress.⁶⁰

If its executives are taken at their word, their company's values sit uncomfortably with the core values of the NHS.

Risk 7

FAILED PILOTS

There is a serious lack of transparency around the data or documentation used to support claims that uses of Foundry have so far have been a success, and NHS professionals are telling us off-the-record that significant problems are being glossed over.

To defend the FDP plan, Palantir and some officials have leant heavily on two supposedly successful uses of the Foundry software – in supporting the vaccine rollout, and in a pilot of the software in a single digitally immature NHS trust, Chelsea and Westminster.⁶¹ These isolated examples, they say, provide sufficient evidence that the £480m FDP is sound policy.

But neither ‘success story’ has opened its data or documentation to scrutiny. And DAUK and Foxglove have consulted multiple NHS managers, data scientists, and health tech experts, all of whom have presented contrary views.

These experts say that the FDP system is poorly designed; that far from delivering on its promises, several Foundry pilots have failed, though this has never been acknowledged; that the FDP unnecessarily duplicates existing NHS infrastructure, which could be developed and iterated for a fraction of the cost; and that the overall plan to have a single private US provider to manage NHS data infrastructure is far too dangerous and expensive.⁶²

From confidential sources we heard that further pilots of the ‘elective care recovery system’ have been indefinitely suspended. These pilots were trialling one

of the envisaged use cases of the FDP. An answer from the DHSC to a parliamentary written question by David Davis MP confirmed this was the case, and that a further nine are currently “paused”.⁶³ The government has not elaborated on the distinction between a “suspended” and a “paused” pilot, nor on the reasons for these suspensions and pauses. But we are aware from NHS sources that several of these suspensions reflect failure of the system by the judgment of the NHS clinical staff responsible for testing them.

There’s surely something not quite right about parliamentarians having to tease information about these pilots via minimalist answers to precisely worded written questions. And the information which has come to light suggests a rather different picture than that which is being pushed by backers of the NHS Palantir deal.

Without transparency about the fate of all Palantir pilots, it is impossible to understand what went wrong, to learn from these failures and assess the overall value for money the FDP tender represents. This presents obvious risks to the soundness of the decision-making process.

ALTERNATIVES

Superior alternatives lie within the NHS's grasp. None of the Palantir FDP sceptics within the NHS lack belief in the potential of data to significantly improve the NHS. But they do have grave doubts about the handling of this procurement and the choice of Palantir as a NHS partner. They also think the FDP gives short shrift to the NHS's own capacity, and the possibility of building on existing solutions

Several of those we spoke to reported there has been an active steer from senior leadership away from the NHS developing its own in-house solutions, or drawing on open-source projects. Yet there are recent, extremely encouraging examples of successful NHS projects which have been developed along these lines.

OpenSAFELY is a Trusted Research Environment developed through collaboration between the Bennett Institute for Applied Data Science at the University of Oxford, the LSHTM EHR research group, TPP and EMIS, and NHS England/NHSX. It's secure, open-source, transparent, and has delivered major contributions to public health during the COVID-19 pandemic. Although designed for research, openSafely's architecture is a national model that is worth expanding and replicating.⁶⁴ But instead it is reportedly struggling to secure onward funding – perhaps because the £480m to Palantir would usurp its functionality.⁶⁵

OneLondon has created a shared data environment for healthcare services covering some 9 million Londoners, which underpins the London Care Record which

enables a range of health and care staff to access the latest details about a patient when they need it.⁶⁶

The new study from Imperial College London also highlights that these are just two of many examples. While there is certainly a job to do in tidying up flows and improving consent, as the study shows, there are extensive data flows and data assets in the NHS. In this sense Palantir's system will largely duplicate – sit on top of – data flows which already exist.⁶⁷

Even within the flawed, over-centralised framework of the FDP, a UK-grown consortium put in an alternative bid. The consortium included companies such as Voror Health Technologies, Eclipse and Black Pear with a much more extensive track record of working with health data, and with the NHS, than Palantir.⁶⁸ This consortium fell at the first hurdle⁶⁹ – perhaps inevitably given Palantir's incumbent advantage, well-resourced lobbying operation, and well-placed cheerleaders within government.

None of the Palantir FDP sceptics within the NHS lack belief in the potential of data to significantly improve the NHS. But they do have grave doubts about the handling of this procurement, and the choice of Palantir as a NHS partner.

RECOMMENDATIONS FOR POLICYMAKERS AND LEGISLATORS

This report has listed several signs that the FDP programme is going off track. The scope is vague. The procurement process is fatally flawed. A strategy of papering over public concerns about privacy, and about Palantir, looks unlikely to hold in the long term.

There's a high level of apparent evasiveness about how the pilot programmes are going. If the FDP does limp on to completion, there's a high chance the NHS will be hindered by the platform's over-centralisation and lack of interoperability.

Ideally, the programme's leadership would recognise these warning signs, pause and reset. However, support by some central managers appears very entrenched, perhaps a testament to the effectiveness of Palantir's lobbying operation.

It therefore falls to our elected politicians to step in. We recommend the following course of action to parliamentarians:

Parliament should investigate the FDP procurement: its design, its value for money, its fairness and the risks the system, as proposed, poses to patient trust.

The FDP tender should be paused until these questions are answered. If no satisfactory answers are available, the tender should be withdrawn.

The government should consult on and design alternatives for managing patient data that better use existing NHS capacity. Where external support is required, any new system should avoid locking in the NHS into one monopoly provider.

The government should reform and clarify how patients can consent to (or opt out of) data-sharing, giving patients more say in how their health record is used.

CONCLUSION

—
It's not too late. There are alternatives - many of which would build on existing projects and in-house NHS expertise.

Improved integration and analysis of health data presents some huge opportunities for the NHS. If it gets its approach right, it could drive improvements both in day-to-day delivery and operational efficiency, and in developing treatments of the future. But the past is littered with failed big ticket NHS data projects, and right now the FDP is doubling down on past mistakes.

Unless there's a course correction, there's a real risk that the FDP becomes either yet another NHS data project that collapses under the weight of public opposition and patient opt-outs, or becomes an albatross round the neck of the NHS, failing to deliver the NHS's core requirements, draining cash, and stifling innovation.

It's not too late. There are alternatives – many of which would build on existing projects and in-house NHS expertise. We hope this report helps persuade politicians and NHS leaders to hit the pause button on Palantir whilst those alternatives are properly explored.

Endnotes

- 1 https://www.theregister.com/2023/02/22/nhs_tech_job_vacancies/
- 2 All figures, unless otherwise stated, are from YouGov Plc. Total sample size was 1,756 adults. Fieldwork was undertaken between 20th - 21st May 2023. The survey was carried out online. The figures have been weighted and are representative of all England adults.
- 3 <https://www.nao.org.uk/wp-content/uploads/2023/03/012229-BOOK-Lessons-learned.pdf>
- 4 The figure is derived from an EY report:
https://assets.ey.com/content/dam/ey-sites/ey-com/en_gl/topics/life-sciences/life-sciences-pdfs/ey-value-of-health-care-data-v20-final.pdf. See also <https://www.thetimes.co.uk/article/us-tech-firms-want-access-to-10bn-nhs-health-data-zpqwkj6pp>
- 5 <https://www.bbc.com/news/health-26259101>; <https://www.wired.co.uk/article/a-simple-guide-to-care-data>
- 6 <https://www.wired.co.uk/article/care-data-nhs-england-closed/> This move followed the publication of National Data Guardian Dame Fiona Caldicott's report on NHS patient data security: <https://www.gov.uk/government/speeches/review-of-health-and-care-data-security-and-consent>. Dame Fiona's review is here: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/535024/data-security-review.PDF
- 7 <https://www.theguardian.com/society/2021/aug/22/nhs-data-grab-on-hold-as-millions-opt-out>
- 8 <https://www.ft.com/content/a9228ce5-8662-474a-9b77-ca563cc5080a>
- 9 <https://www.digitalhealth.net/2022/12/federated-data-platform-palantir-juggernaut-continues/>
- 10 <https://www.theguardian.com/society/2022/nov/13/controversial-360m-nhs-england-data-platform-lined-up-for-trump-backers-firm>
- 11 <https://www.theguardian.com/politics/2020/sep/17/uk-awards-border-contract-to-firm-criticised-over-role-in-us-deportations>
- 12 <https://nymag.com/intelligencer/2020/09/inside-palantir-technologies-peter-thiel-alex-karp.html>
- 13 <https://www.palantir.com/offerings/health/>
- 14 https://www.theregister.com/2023/01/04/palantirs_covidera_uk_health_contract
- 15 <https://www.politico.eu/article/uk-trade-department-courted-us-firm-palantir-over-untapped-nhs-data-potential/>
- 16 <https://www.thebureauinvestigates.com/stories/2021-02-24/revealed-data-giant-given-emergency-covid-contract-had-been-wooing-nhs-for-months> Prior followed up to thank Mosley, saying "Thank you for hosting such an interesting dinner and also for the water melon [sic] cocktails! If you can see ways where you could help us structure and curate our data so that it helps us deliver better care and provides a more insightful data base for medical research do be in touch."
- 17 Lord Prior has recently resurfaced in the news for lobbying NHS officials on behalf of other private US health firms: <https://www.theguardian.com/society/2023/feb/19/ex-nhs-england-chief-david-prior-teladoc-lazard>
- 18 After openDemocracy and Foxglove forced the contract to be published, it was clear the original contract kept all of the project IP with Palantir – in other words, it let the firm refine its tools from its emergency access to patient data. <https://www.opendemocracy.net/en/ournhs/under-pressure-uk-government-releases-nhs-Covid-data-deals-big-tech/>
- 19 <https://healthtech.blog.gov.uk/2020/03/28/the-power-of-data-in-a-pandemic/>
- 20 See <https://www.nytimes.com/interactive/2020/12/17/world/europe/britain-Covid-contracts.html>; <https://www.bmj.com/content/375/bmj.n2825>
- 21 https://www.theregister.com/2023/01/04/palantirs_Covidera_uk_health_contract/
- 22 <https://www.washingtonpost.com/business/on-small-business/peter-thiels-palantir-had-secret-plan-to-crack-uks-nhs-buying-our-way-in/2022/09/30/c62029a8-4081-11ed-8c6e->

- 9386bd7cd826_story.html
- 23 <https://healthtech.blog.gov.uk/2020/03/28/the-power-of-data-in-a-pandemic/>
- 24 <https://www.digitalhealth.net/2022/06/second-nhs-official-to-join-palantir-as-it-guns-for-360m-contract/>
- 25 <https://www.ft.com/content/3f6f24f8-9e5c-42c3-8ae6-bfef5f953524>
- 26 For example, ‘purpose-based access controls’.
- 27 <https://www.ft.com/content/11f250df-5125-4bab-bbee-3c4324685c55>
- 28 <https://www.nao.org.uk/wp-content/uploads/2020/09/Lessons-learned-from-Major-Programmes.pdf>
- 29 These are “population health and person insight; vaccination and immunisation; elective recovery; care coordination; and supply chain.”
- 30 Document C104091_NHSE_FDPAS_Document1_Prospectus.pdf – p.11 <https://health-family-contract-search.secure.force.com/servlet/servlet.FileDownload?file=00P8d000009dbn6EAA>
- 31 Ibid. at 20.
- 32 <https://www.digitalhealth.net/2022/12/federated-data-platform-palantir-juggernaut-continues/>
- 33 <https://www.digitalhealth.net/2023/05/nhs-england-says-all-nhs-trusts-to-get-own-fdp/>
- 34 C104091_NHSE_FDPAS_Document2_IG_CS.pdf – p. 5 <https://health-family-contract-search.secure.force.com/servlet/servlet.FileDownload?file=00P8d000009db-dIEAA>
- 35 Document C104091_NHSE_FDPAS_Document1_Prospectus.pdf (supra note 27) at 40.
- 36 <https://www.nao.org.uk/wp-content/uploads/2019/06/Challenges-in-using-data-across-government.pdf>
- 37 <https://www.theguardian.com/society/2021/aug/22/nhs-data-grab-on-hold-as-millions-opt-out>
- 38 <https://digital.nhs.uk/data-and-information/publications/statistical/national-data-opt-out/january-2023>
- 39 A failure to account for these concerns has caused the failure of planned patient data centralisations for years, from the National Programme for IT to care.data to 2021’s debacle, the General Practice Data for Planning and Research, which saw over a million new NHS patients opt out of sharing data within a month.
- 40 <https://www.independent.co.uk/news/uk/politics/nhs-pressured-reveal-confidential-patient-data-home-office-ministers-theresa-may-secretary-digital-kingsley-manning-a7558136.html>
- 41 It is also worrying that in the place of the independent government safe haven for data – NHS Digital, now abolished – there will stand only Palantir, and the same officials in NHS England who have designed the FDP.
- 42 Document C104091_NHSE_FDPAS_Document2_IG_CS.pdf, pp.4, <https://health-family-contract-search.secure.force.com/servlet/servlet.FileDownload?file=00P-8d000009dbdIEAA>
- 43 For example, people may want to opt out of commercial uses (not to share with eg Google) but be eager to share data for secondary planning within the NHS. If they engage the National Data Opt Out, both are gone.
- 44 https://datainsights.uk/wp-content/uploads/2023/03/data_flow_manuscript_1.0.pdf
The Imperial study also recommends reforming the opt-out to let people decide who accesses the data: “Opt-out conditions should be set at the level of distribution to types of consumers, rather than at extraction. As capabilities advance, binary extraction opt-outs will limit patient autonomy and restrict access to key data-driven interventions.”
- 45 https://www.theregister.com/2023/01/11/nhs_england_launches_procurement_for/
- 46 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/774071/MoU_between_HSCIC__Home_Office_and_DH.pdf; see NHS

Digital's Statement on the Memorandum at <https://digital.nhs.uk/news/2018/a-statement-from-nhs-digital-on-the-memorandum-of-understanding-with-the-home-office>

47 https://datainsights.uk/wp-content/uploads/2023/03/data_flow_manuscript_1.0.pdf
 “New data infrastructure must focus on expanding extraction layer flows, rather than reshuffling existing assets into additional nodes of dissemination. This may require technologically individualised solutions across regions. Most pertinently, while a new national Federated Data Platform (FDP) will enable analytics across regional data environments that are held by integrated care systems, it is as yet unclear how the resulting data asset will differ from population data already held by the NHS, beyond transferring control over data aggregation to a different party.”

48 <https://www.onelondon.online/about/>

49 <https://echalliance.com/better-wins-contract-to-implement-shared-care-plan-solution-for-one-london/>

50 <https://www.england.nhs.uk/blog/better-insights-better-decisions-better-health/>

51 The NHS' decision to 'split' the privacy and other modest parts of the FDP contract does not go far enough to assuage these concerns; the core infrastructure is still set to be provided by one supplier for one vast national database.

52 This is also consistent with earlier reporting of US police departments' problematic experience with Foundry, which found themselves locked into a system with escalating costs and poor interoperability: <https://www.wired.com/story/how-peter-thiels-secretive-data-company-pushed-into-policing/>

53 <https://www.palantir.com/newsroom/letters/letter-to-shareholders/august-8-2022/en/>

54 <https://www.bbc.co.uk/news/uk-politics-55978334>

55 Kingsley Manning [the former head of NHS Digital] wrote about this in the BMJ, “Doing Away with NHS Digital: Why It Matters”, at <https://www.bmj.com/content/376/bmj.o361/rr-0>

56 At an on-record meeting with Foxglove staff, we asked Palantir's Louis Mosley directly what he would do were UK government officials to seek to use health data for collateral purposes – such as border enforcement. Palantir's staff on the call gave us a partial answer, referring to the need safeguards – but we understood the upshot of what they said was that if government officials came for the data, Palantir would not stand in their way. Foxglove's contemporaneous note of this conversation is on file and we would be happy to share it with interested officials.

57 <https://www.euractiv.com/section/artificial-intelligence/news/german-constitutional-court-strikes-down-predictive-algorithms-for-policing/>. Several UK police departments have refused to answer FOI requests querying whether they also use Palantir software.

58 <https://www.bloomberg.com/news/articles/2023-01-18/peter-thiel-says-british-affectation-for-nhs-is-stockholm-syndrome?accessToken=eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJzb3VyY2UiOiJTdWJzY3JpYmVpY2ImdGVkQXJ0aWNsZSIsImh0dCI6ImTY3NDA1OTk5NCwiZXhwIjoxNjc0NjY0Nzk0LCJhcnRpY2xiSWQ0IjST09QU1ZEV0xVN-jgwMSIsImJjb25uZWN0SWQiOiIzNjREMjRGNDQ0OTEyQjE0Mzc1OTA4ODY4N0FF-NiJ9.fn7f1TD2apbb2sTU5D7NQu2KH3BwJHlewR-5VOGUTJA>

59 <https://washingtontechnology.com/2021/02/palantir-details-new-phase-of-federal-strategy/355117/>

60 <https://medium.com/@ssankar/innovation-needs-customers-not-capital-e18b8e55705a>

61 We encourage Parliament, in assessing the FDP, to secure documentation of these 'success cases' as well. These uses of Foundry may well have worked – but without scrutiny these stories are not data science. They are marketing.

62 This £480m tender, colossal as it is, is the first of four envisaged data management procurements.

63 <https://members.parliament.uk/member/373/writtenquestions#expand-1601903>

64 Discussion of the openSafely platform is contained in the Imperial study mapping ex-

- isting NHS data flows, see p21: https://datainsights.uk/wp-content/uploads/2023/03/data_flow_manuscript_1.0.pdf
- 65 <https://www.opensafely.org/>
- 66 <https://www.onelondon.online/about/>
- 67 https://datainsights.uk/wp-content/uploads/2023/03/data_flow_manuscript_1.0.pdf
- 68 <https://www.ft.com/content/11f250df-5125-4bab-bbee-3c4324685c55>
- 69 <https://www.hsj.co.uk/technology-and-innovation/nhse-rejects-british-alternative-bid-to-run-480m-data-platform/7034610.article>

Acknowledgements

This report draws on extensive input from a huge range of NHS professionals - we're very grateful for their time, and their courage, in speaking out to us.

Respecting their wishes these sources are unnamed. Several have however indicated they are willing to speak to legislators.

None of these people necessarily agree with any of the report's conclusions and any mistakes are the responsibility of DAUK and Foxglove alone.

This report was written by Cori Crider for Foxglove, with support from David Babbs and Tom Hegarty.



—
**THE NHS FEDERATED DATA
PLATFORM AND PALANTIR**

—
7 KEY RISKS

—
June 2023