

Ten years of *innovation*. Tens of millions in *savings*.

Why the access layer, not the application layer, is where UK government can move fastest

Public Sector Point of View

Network & Access Edition

July 2026

The opportunity in front of CDIOs

Departmental technology leaders have been handed something they have asked for repeatedly and rarely received: clear authority to change things, and an explicit mandate to do it for less money.

The machinery around them has been rearranged again. Accountability for digital, data and artificial intelligence has moved. Central spend controls have given way to delegated departmental authority. Fiscal headroom has narrowed, and the next Budget will reward savings that can be demonstrated inside the current spending period rather than promised beyond it. The specifics will keep moving. Three consequences hold regardless.

Decision rights have moved closer to the department. A well-scoped technology decision can now be taken and approved departmentally rather than negotiated with the centre. That materially changes what a CDIO can commit to, and how fast.

Cost is being scored on a shorter clock. A modest saving that lands this year is worth more, politically and operationally, than a larger one that lands after the next election. The currency has changed from ambition to evidence.

The expectation on AI has become concrete. Departments are being asked to demonstrate adoption rather than intent, and to demonstrate it with controls already in place.

Taken together, this is the most permissive environment for departmental technology decisions in a decade. It is also the least forgiving of any programme that needs five years to produce its first benefit.



The opportunity is real. The risk is that departments answer a short-clock mandate with a long-cycle instrument.

The trap inside the mandate

Most CDIOs already know what their estate needs, and the published departmental roadmaps are, on the whole, right. **The constraint is not diagnosis. It is the delivery model the sector reaches for by default.**

Legacy remediation in government has historically meant one shape of intervention: a large replatforming programme, running five to ten years, with high capital cost, significant operational disruption, and a benefits case landing well beyond the horizon of the administration that approved it. The record on this is not ambiguous. Twenty-eight per cent of central government IT still runs on legacy technology, up from 26% in 2023. On the most recent published analysis, around £2.3bn a year, close to half of central government IT spend, goes on keeping it alive. The National Audit Office has costed roughly £3bn in delays and resets across five major digital change programmes, and found departments still locked into a small number of large suppliers with substantial awards made without competition.

A department that commissions a replatform this year books the disruption now and the benefit in the next parliament. That is not a failure of ambition. It is a mismatch between the instrument and the mandate, and it is the single most expensive mistake available to a CDIO over the next twelve months.

Reorganisation changes none of it. **Machinery of government changes move accountability. They do not move code.** Every department that carried a red-rated legacy estate last month carries the same one today, along with the same device fleet, the same supplier lock-in and the same technical debt.

One thing has genuinely changed in the department's favour, and it is commercial. Central digital spend controls have been replaced by delegated departmental authority under the Digital Assurance Playbook. Set alongside the Procurement Act 2023 competitive flexible procedures and a stated preference for contracting directly with smaller suppliers, **the approval path for a well-scoped technology decision is shorter today than at any point in the last ten years.**

Which makes the question worth asking a narrow and answerable one:

What can a department change inside a single financial year, without touching its applications, that measurably improves the experience of using them, reduces the security exposure around them, and takes real cost out of the run?

The insight: change the control point, not the application

There is one move that lets a CDIO deliver innovation and savings on the same clock, and it is deceptively simple. Stop trying to modernise the application estate first. Modernise the layer through which every application is reached.

The great majority of government work now reaches the user through a browser. GOV.UK services, Microsoft 365 and Copilot, contact centre platforms, CRM, case management, and the internal web front ends bolted onto systems older than the civil servants using them.

The browser is the one place where all of it converges. **In most departments it is also the only part of the stack that nobody governs.** Policy is enforced on the device beneath it and the network around it, at considerable expense, while the layer where citizen data is actually handled is a consumer product downloaded from the internet.

An enterprise browser turns that layer into a control plane. Same applications. Same back ends. Same migration plan, unchanged. Different control point.

The consequence is the part that matters to a business case. **Nothing in the estate has to be rewritten for this to work, which is exactly why it can be delivered in months rather than years.** A department is not modernising the application. It is modernising how people reach it, what they are able to do with the data once inside it, and what it costs to give them access at all.

This is what makes the mandate deliverable rather than contradictory. The replatforming still needs to happen and should still be planned. **Putting a modern, governed access layer over the estate first gives a department the user experience, the security posture and the cost base of a modernised organisation while the modernisation is still in progress.** Five benefits follow directly, and they arrive together rather than in sequence.

1. Earlier financial benefit: the third-party device estate stops scaling with demand

Be clear about where the money actually is, because the obvious answer is the wrong one. Civil servants will keep their managed laptops, and they should. Many are lower-paid workers who cannot reasonably be

expected to supply their own equipment, and a department that tried to move that cost onto its own staff would deserve the coverage it got. Departmental employment is not a bring-your-own-device proposition, and nothing here argues that it should be.

The avoidable cost sits with everyone else. **Government policy is corporate by default: anyone touching corporate data gets a managed device.** The policy itself acknowledges that not all of its requirements automatically suit third-party suppliers, and asks departments either to write bespoke contractual security requirements or to issue a corporate device. Most issue the device, because the alternative is hard to evidence.

So departments ship laptops to contractors, outsourced contact centre agents, system integrator staff, seasonal surge capacity and supplier teams who already own perfectly serviceable equipment. **Every one of those devices carries a capital cost, an imaging and logistics cost, a management licence, a service desk overhead and a recovery problem at the end of the engagement.** That is the line worth attacking, and it is where this argument is strongest.

Two further reductions follow on the internal estate without taking anyone's laptop away. Virtual desktop seats issued alongside a laptop, purely to reach a handful of applications, stop being necessary. And where an endpoint's job narrows to hosting a governed browser, **device specifications can come down and refresh cycles can stretch.** That one shows up quietly across a large fleet rather than dramatically in any single year, and it should be claimed modestly.

The estate divides into three groups, and all three end up in the same window.

- **Web applications only:** The largest group by some distance. Microsoft 365, GOV.UK services, CRM, case management, the contact centre platform. A governed browser is a straight substitution, and it improves on what these users have today, because the controls follow the person rather than the hardware.
- **Complex internal applications:** Older line-of-business systems, internal portals and legacy web front ends that currently need a

corporate device, a private network path or a separate client. These reach the user through the same browser, in one window, behind one sign-in. That matters as much for experience as for cost. Staff today move between a desktop, a remote session and a browser to finish a single task, and they stop having to.

- **Applications that genuinely need a virtual desktop:** A much smaller share than most departments assume. Even these can be published as individual applications inside the browser rather than as a full remote desktop session, so the department pays for the application instead of a second Windows desktop per person, and the user never leaves the window they were already in.

The timing is unusually good. A number of large managed desktop and workplace services contracts across government reach expiry and re-competition during 2027. Decisions taken in the next twelve months will set the cost base for the seven years after that.

2. A better employee experience, and it is measurable

This is the benefit most often left out of the business case, and it is the one civil servants will actually notice. Access friction is a daily tax. It shows up in engagement scores, in time to productivity, and in the volume of service desk tickets that exist only because getting to a system is hard.

Four things change on day one for the people affected:

Onboarding drops from weeks to the same day

A new starter or a supplier's agent installs a browser, signs in through the identity provider, and reaches exactly their permitted applications. No build, no shipment, no desk visit.

Signing in stops being a wait

Users of the outsourcing platform yoummday went from around two minutes to log in through a virtual desktop to about ten seconds through the browser. At a contact centre with thousands of shift starts a day, that is a schedule change, not a nicety.

Offboarding becomes instant and provable

Access is withdrawn at the identity provider and there is nothing on the endpoint to recover. Compare that with a device that has to be chased, may never come back, and remains a data risk on somebody's kitchen table until it does.

One window instead of three

Web applications, complex internal systems and published remote applications all appear in the same place, behind one sign-in, and it behaves like the browser people already use. Nothing to relearn, and the assistive technology that works today keeps working, because the rendering and extension model is the one it already supports.

All of this is measurable, and a department should insist on measuring it. Baseline time to first productive day, service desk volume for access issues, and user satisfaction with the current device and virtual desktop experience before the pilot starts. Those three numbers make the strongest part of the case, and they are the part a minister can repeat.

3. Operational resilience: the extended workforce becomes securable in a day

Departments are expanding their use of managed service providers, outsourced contact centres and system integrators at speed, and much of that workforce touches citizen data. Under corporate by default, every one of those people needs a shipped laptop or a virtual desktop seat.

That model is already failing in practice. Onboarding runs into weeks. Devices go out and do not come back. Peak demand cannot be met through hardware logistics, because you cannot ship ten thousand laptops for a filing deadline and collect them afterwards.

Browser-delivered access inverts it. A third-party agent signs in on their own device, posture checks apply, and they reach exactly the applications their role permits. Copy, paste, download, print and screenshot are

governed at the point of use. Every session is watermarked and audited. Deprovisioning is instant through the identity provider, because nothing sits on the endpoint to recover.

The resilience gain is broader than the extended workforce, and it tends to reveal itself under stress. Surge capacity becomes a licensing question rather than a logistics one. Access no longer depends on a specific device, a specific network path or a specific building. The University of the Pacific reports that early in its deployment, during an outage caused by another supplier which took down a wide range of its internal systems, its team used Island to reach the internal network, fix the problem and restore service. They had not planned that test.

4. Reduced supplier complexity: point tools collapse into one platform

Most departmental security estates have accumulated separate products for data loss prevention, remote browser isolation, secure web gateway, privileged access, virtual desktop and, increasingly, AI usage governance. Each carries its own licence, its own integration debt and its own operational team.

Every one of those controls is, functionally, a control over what happens in a browser session. Moving them into the browser retires products rather than adding one. Departments whose published strategies commit them to fewer, better-supported platforms are describing this exact outcome, and it is available as a product rather than a programme.

5. Faster AI adoption: governance arrives at the same time as the AI

Copilot rollouts across UK government already run to tens of thousands of seats, and the Cabinet Office now owns public sector AI adoption directly. The controls have not kept pace with the licences.

In-browser governance is the only place to apply them without slowing adoption down. Control which AI tools staff can reach. Inspect and redact what is pasted into a prompt before it leaves. Log every interaction for audit. Block the unsanctioned ones. This is deliverable in weeks, and it is the difference between an AI programme that survives its first NAO review and one that does not.

The network decision, and how to avoid buying new tech debt

A second decision is arriving at the same time as the first, and it will lock in an architecture for the rest of the decade.

Departments are buying secure access service edge now. The Crown Commercial Service Network Services 3 framework, which covers SASE under Lot 1a, expires on 17 July 2027, and several large managed workplace contracts land in the same window. **Two of the most consequential architectural decisions a department will take this decade are being made months apart, by different teams, usually without reference to one another.**

SASE was defined in 2019 and its premise was sound. Users were leaving the office, applications were leaving the data centre, and the inspection point had to move from an appliance in a building to a service in the cloud. As a replacement for backhauling every branch through a corporate firewall, it was a clear improvement.

The premise has aged in four specific ways, and each one bites harder in government than in the average enterprise.

The risk moved inside the session

An outsourced agent entitled to open a citizen record, who then copies it somewhere it should not go, generates traffic indistinguishable from doing the job properly. Network inspection sees an authenticated user reaching a permitted application over an encrypted connection. **That is the entire event, and it is the shape of most real data loss in a service organisation.**

Inspection is closing

TLS 1.3 removed the cleartext metadata that made selective inspection cheap. QUIC encrypts almost the whole handshake. Encrypted Client Hello hides the destination hostname from a proxy altogether, and the mainstream industry answer is to detect Encrypted Client Hello and block it. **The architecture now depends on preventing a privacy feature from working.** Microsoft's own Global Secure Access does not support TLS 1.3 with Encrypted Client Hello, because an encrypted server name leaves nothing to build a certificate against. AWS Network Firewall drops the connection. Each year the share of traffic a man in the middle can usefully read gets smaller, and the cost of pretending otherwise gets larger.

These are primary-source claims and should be checked rather than taken on our word. Microsoft Learn, Global Secure Access transport layer security FAQ. AWS Network Firewall developer guide, TLS inspection considerations. IETF draft-camplig-ech-deployment-considerations. All three are public.

The AI problem is not a traffic problem

What matters is the content of a prompt box before anyone presses send, and whether it contained a National Insurance number. That is a decision about a rendered page, taken at the keyboard. Reconstructing it from a packet stream is difficult, expensive and late.

The population that needs it most cannot run it

SASE reaches a device with an agent installed. The contractors, outsourced agents and supplier staff who carry the sharpest data risk work on equipment the department does not own, cannot instrument and has no right to manage.

The control does not reach the exposure.

Why signing now creates the debt

A SASE commitment at departmental scale runs three to five years, ties the department to one vendor's point-of-presence footprint and one agent on every endpoint, requires a standing operations team, and becomes the reference architecture every later access decision has to accommodate. None of that is unusual. None of it is wrong in itself.



Tech debt is not created by keeping old things. It is created by signing long contracts for the current generation of things, at the moment the generation turns over.

A five-year commitment to network-centric inspection, signed in the year the protocols are closing that inspection off, delivers a department into 2031 holding a modern-looking contract and a dated architecture. It will have paid for the transition and arrived where it started.

Inspect by exception, not by default

The answer is not a smaller SASE contract. It is a different default.

Today's model is to route all user traffic somewhere for inspection, then permit it. That made sense when the browser could not be trusted to enforce anything. It no longer holds. **Enforce at the point of use, in the session, and route traffic out for inspection only in the specific cases that genuinely require it.** For the great majority of user activity, the browser is where the decision belongs, because it is the only place the rendered page, the user's action and the data are all visible at once.

The residual is real but small: non-browser and thick-client applications, server-to-server traffic, DNS-layer controls, operational technology, and unmanaged equipment on a departmental network. **A modern enterprise browser platform provides that network capability itself, rather than requiring a separate SASE estate alongside it.** Island delivers modern SASE as part of the platform, with inspection applied by exception rather than to everything by default. That is a materially different commercial and architectural position from buying a browser and a network contract and paying twice for overlapping control.



Control	Legacy SASE and SSE	Browser-first, inspect by exception
Default posture	Inspect all traffic, then permit	Enforce in the session; inspect by exception
Where policy applies	Between device and destination	Inside the authenticated session
Unmanaged and third-party devices	Agent required, often not possible	Native
Encrypted traffic	Interception, narrowing each year	Reads the rendered page, no interception
Copy, download, screenshot in a permitted app	Little visibility	Governed action by action
Content of an AI prompt	Difficult and late	Native
Non-browser and server traffic	Core strength	Same platform, applied by exception
Typical commitment	Three to five years, agent and PoP tied	Software, per user

Sequence it the other way round

Put the browser control point in first. It deploys in weeks, needs no agent on equipment the department does not own, and covers the population and the activity carrying most of the data risk. **Then measure what is actually left.**

Our claim is that the genuinely non-browser residual is smaller than most incumbent network proposals assume. **That is a hypothesis, and it is one a department can settle for itself.** The flow data needed to test it already exists inside the network team, and a fortnight of analysis will establish the real split between browser-borne user traffic and everything else. **Run that analysis before committing to either architecture.** If the residual turns out to be larger than we expect, the sequencing argument still holds and the network scope is simply bigger than we assumed.

Buy in that order and the department commits to inspection capacity for the exceptions rather than for the whole estate, on a shorter contract, with a working alternative already in production. Buy in the other order and it signs for full scope, deploys it, then discovers it still needs a control point inside the session. **At that point it is paying twice, and the architecture is fixed until the contract ends.**

The future-proofing argument is the simplest one. A control point in the session survives changes in network topology. **A control point in the network does not survive changes in what applications do.** Whatever arrives next, and browser-based AI agents acting on a user's behalf are already here, will appear inside a session first and show up on the network as ordinary encrypted traffic to a permitted destination.



Speed to value, in one financial year

Nothing in this depends on a migration finishing. Each stage stands on its own, delivers a measurable result, and can be stopped without stranding the one before it.

Weeks 1 to 4: Pilot

One persona, one set of applications. Policy baseline set. Data controls, session recording and AI guardrails live for the pilot group. No infrastructure change, no application change.

Months 2 to 3: Extended workforce

Contractors and outsourced agents onboarded on unmanaged devices. Time to access drops to the same day. First point product identified for retirement. First audit evidence produced.

Months 4 to 6: Browser-first personas

Task workers move off managed laptops and virtual desktops. Complex internal systems delivered in the same window. Remote applications published as apps, not desktops. Savings begin landing in-year.

Months 6 to 12: Estate and procurement

Rollout across the estate. Workplace and network re-procurements scoped to the residual, with a working alternative already in production. Full run-rate saving carried into the next financial year.

First controls live in **weeks**. First savings inside **the current financial year**. Full run-rate by **month twelve**. The comparison worth putting in front of a minister is that an application replatform started in the same month produces its first measurable saving in **2031**.

Deployment speed is not a theoretical claim. Mattress Firm rolled Island out to roughly 2,400 stores and 7,000 devices in under two weeks. TaskUs deployed to 35,000 people across 300 segmented business units and reported no downtime and no service interruption during the rollout. Both estates are more fragmented than most single UK government departments, and TaskUs is considerably larger.

The arithmetic, at government scale

These are potential savings, deliberately modelled on the conservative side, and they are not committed or guaranteed. We have taken no credit for removing managed laptops from civil servants, because that is neither realistic nor desirable. The model assumes internal staff keep their devices. What it counts is the third-party estate, the virtual desktop seats issued on top of a laptop, a modest specification and refresh effect, and the security products the browser makes redundant.

Assumptions, large department

- 80,000 internal staff, all of whom keep a managed laptop, plus 10,000 third parties, contractors and supplier staff currently issued departmental hardware or a virtual desktop seat
- Fully loaded managed device at £950 per user per year, covering four-year amortisation, imaging, logistics, management licence and service desk. Virtual desktop seat at £1,400 per year, all in
- Third-party provisioning moves to the individual's own equipment, reducing that cost by 50% to 65% rather than eliminating it
- 5,000 internal virtual desktop seats issued alongside a laptop, of which the browser displaces 50% to 60%
- 30,000 internal endpoints eligible for lower specification and longer refresh, at a cautious £100 to £150 per user per year
- Network inspection scoped to the exceptions at the 2027 re-procurement rather than to all user traffic



Potential high	Basis	Line	Potential low
Third parties provisioned without departmental hardware	10,000 users, £1,050 blended baseline	£5.3m	£6.8m
Virtual desktop seats retired from the internal estate	5,000 seats, £1,400 baseline	£3.5m	£4.2m
Device specification and refresh cycle effect	30,000 endpoints, £100 to £150	£3.0m	£4.5m
Point security products consolidated into the browser	80,000 users, £30 to £45 blended	£2.4m	£3.6m
Network inspection applied by exception, not by default	80,000 users, £15 to £30 avoided	£1.2m	£2.4m
Onboarding and offboarding time recovered	Weeks to same day, per joiner	n/a	n/a
Total	Potential annual run-rate saving	£15.4m	£21.5m

Potential savings only. An illustrative model, not a quotation and not a commitment. Baselines are drawn from published public sector end-user compute benchmarks. The third-party line is avoided cost as that workforce scales rather than a cash reduction against today's budget, and should be presented separately in a business case. The refresh line accrues over a cycle rather than in year one. The network line is avoided future commitment at re-procurement, not a saving against a contract already signed.

Call it a potential £15m to £21m a year in a large department, or £77m to £108m across a five-year term, without a line of application code being rewritten. That is a smaller number than a vendor would usually put in front of you, and it is the one we would defend in front of a finance director.

Scale it down honestly. **A mid-sized department of 20,000 staff with 2,500 third parties models at roughly £3.8m to £5.4m a year, or £19m to £27m over five years.** The mechanism does not change with size. The absolute number obviously does.

Expressed as a unit rate, the model runs at **roughly £190 to £270 per member of staff per year.** That is the number worth carrying into a conversation, because it can be multiplied by any department's headcount without inheriting our assumptions about anyone's estate.



Treat this model as a template to rebuild, not a number to quote. Every figure in it should be replaced with the department's own device, virtual desktop, network and licence actuals before it goes near a business case.

The security case runs alongside it and is harder to price. Stolen credentials remain the most common way into a government network, and the browser is where every one of those credentials gets typed. Keeping credentials and citizen data off the endpoint, applying phishing-resistant checks at the point of entry rather than the point of network access, and recording what a session actually did, all address that class of loss directly. For an extended workforce on devices the department does not own, the browser is the only place those controls can be applied at all.

This is not a pilot technology

The reasonable question from any accounting officer is whether this is a new product being tested on the public sector. It is not. The closest analogue to a government department's problem is the business process outsourcing industry, where tens of thousands of third-party staff need governed access to a client's sensitive systems, on devices the client does not own, at short notice and at low cost. That industry has already made this change.



Business process outsourcing

35,000 staff, 300 segmented business units, no downtime

TaskUs staff act as third-party contractors with privileged access to client back-office systems and sensitive data, which is precisely the position a departmental BPO partner occupies. Its security team could not see what 65,000 people were doing inside a consumer browser.

Island now runs across 35,000 employees, heading to 65,000, spanning 300 highly segmented business units, each with different client security requirements. The rollout was completed by a small engineering team with no downtime and no service interruption. Island replaced a data loss prevention endpoint product and a digital experience monitoring product, each costing hundreds of thousands of dollars, and a planned investment in a cloud access security broker was avoided. Satellite delivery centres now operate without on-premises firewalls that had cost between \$60,000 and \$150,000 each.

"The black box of the traditional browser was not giving us the visibility and control that we needed and not protecting the user data to the level our infosec standards required."

Eric Lowe, Director of InfoSec Risk Management, TaskUs

35,000

Staff deployed, on the way to 65,000

0

Downtime or service interruption at rollout

4x

Faster client onboarding

300

Segmented business units supported

yoummday

Outsourced customer experience, Munich

25,000 freelancers on their own laptops, at 40% below the virtual desktop option

yoummday matches clients with around 25,000 registered freelance professionals who work on their own equipment. It tried distributing security software on USB sticks, then moved to virtual desktops. The virtual desktops were secure but expensive, and could not deliver controls the CISO wanted, screenshot protection among them.

Ramping a new client project took about two days on virtual desktops and now takes about four hours. Freelancer sign-in went from around two minutes to about ten seconds. Island is offered to clients at roughly 40% below the virtual desktop option, and about half of all new projects since 2026 have chosen it.

"Before I found Island, I wasn't aware that a browser could provide such good security features."

Boris Bock, CISO, yoummday

40%

Lower cost than the
virtual desktop option

4 hrs

To stand up a client
project, from two days

10 sec

To sign in, from about
two minutes

25,000

Freelancers on
unmanaged devices

The pattern repeats elsewhere in regulated and critical environments. Teleperformance, one of the largest outsourcing groups in the world, uses Island to apply granular security policy across a global workforce. Pfizer uses it to secure a global workforce in a heavily regulated sector. Swiss Life runs bring-your-own-device access in regulated financial services. Landis+Gyr, in energy metering, replaced virtual desktops and cut VPN usage by 80%. Brightline Trains reports that displacing virtual desktop cost meant the platform essentially paid for itself. Omada Health consolidated eight browser variants into one and reports full payback within six months.

One detail from the TaskUs story deserves a departmental reading. TaskUs first encountered Island because a bank, rather than impose the lengthy security requirements TaskUs expected of a financial services client, simply told them to use it. That is a regulated institution using the browser as its supplier assurance mechanism. It is exactly the move available to a department that needs to hold a BPO partner to a standard without auditing their laptop estate.

Island holds **ISO 27001** and **SOC 2 Type II**, both independently audited. In government specifically, it entered the United States federal market in September 2024 and is available to federal departments, agencies and their contractors through Carahsoft, with support for government-furnished equipment, approved personal devices and mission partner data sharing across networks of different classification. It is certified under TX-RAMP at Level 2 and is available to US state and local government through Carahsoft's purchasing vehicles.

Two caveats belong here rather than in a footnote. Island is FedRAMP In Process rather than fully authorised, so the US federal position is evidence of suitability and market entry, not of completed accreditation. And there is no named UK central government deployment at 80,000-seat scale to point to today. The closest published analogues are outsourcing organisations of comparable and greater size, regulated commercial operators, and US public sector. **Any department should ask for a reference of comparable complexity, and should ask directly about Cyber Essentials Plus, alignment to the NCSC cloud security principles, and where the control plane and its data are hosted for UK customers.** Those are reasonable questions and they deserve documented answers before a pilot becomes a commitment.

Customer figures in this section are published by Island in its customer stories, are reported by the customers named, and are not independently audited. Sources are listed at the end.

What this does not do

Any paper of this kind is worth more if it is explicit about its own limits, so here are ours. A CDIO will find these anyway, and it is better to arrive at the conversation having named them.

This is browser-layer control, and on a device the department does not own it is a compensating control rather than an equivalent to a managed endpoint. **It is stronger than that sounds.** The browser defeats locally installed software keystroke loggers by feeding them false key presses, blocks software screen capture tooling from reading protected windows, encrypts its cache, cookies and stored credentials so they cannot be extracted and replayed, detects attempts to inject code into or tamper with the browser process and can end the session when it finds them, and applies different policy according to device posture. What sits outside its reach is narrower than the usual objection assumes, and still real. **A kernel-level compromise or a hardware keystroke logger sits beneath every user-space application, this one included,** and no software prevents someone photographing a screen. Watermarking deters casual capture and supports attribution afterwards rather than preventing it. **This closes the categories of attack that happen at volume, and leaves a residue that is either physical or requires privileged compromise of the endpoint itself.**

The residual network scope is a claim to be tested, not a finding:

Our position is that the volume of genuinely non-browser traffic in a modern department is smaller than most incumbent network proposals assume. That is knowable from a department's own flow data before anything is committed, and it should be measured rather than accepted from either side of the argument.

There is a dependency on the vendor and on upstream Chromium:

Patch cadence follows Chromium releases, and a department is taking a dependency on a supplier's release engineering and CVE turnaround. Policy is administered from a cloud control plane. Both of those are legitimate questions for a due diligence process, including questions about availability history and how enforcement behaves when the control plane cannot be reached.

Integration and change are real work:

Connecting to an existing identity provider, privileged access tooling and logging pipeline takes effort and should be scoped honestly. The claim that there is no retraining refers to the interface, which behaves like Chrome and Edge because it is built on the same engine. It is not a claim that a rollout to tens of thousands of people needs no communications, no service desk preparation and no handling for the awkward minority of workflows that will not move cleanly.

What ten years actually means

Potential high	Application replatform	Access layer first
First measurable saving	Year three to five	Quarter one
Legacy applications touched	All of them	None
Programme risk	Concentrated, irreversible	Reversible, per persona
Third-party workforce secured	After migration	Immediately
AI controls	Deferred to future state	Live at deployment
Business case horizon	Next parliament	This Budget

These are not alternatives. The replatform still has to happen, and it should. Putting a modern, governed access layer over the estate first gives a department the user experience, the security posture and the cost base of a modernised organisation while the modernisation is still under way. It also front-ends the legacy systems that will be last in the queue, adding multi-factor authentication, conditional access, data controls and session recording, without asking anyone to reopen thirty-year-old code.



**That is the leapfrog. Not skipping the work.
Getting the outcome first.**

Where Island fits

Island is the enterprise browser. It is built on Chromium, so it looks and behaves like the browsers people already use every day, Chrome and Edge alike, with the same tabs, the same extensions, the same shortcuts and the same rendering. There is no retraining and very little user resistance, which is the usual reason access projects stall.

Underneath, every function departments currently buy as a separate product becomes a policy setting: data loss prevention, isolation, web filtering, privileged access, session recording, device posture, and governance of what staff can put into an AI tool. Web applications, complex internal applications and published virtual desktop applications all arrive in the same window, so a person doing three different jobs sees one consistent experience rather than three.

For a UK government department that translates into five decisions, each of which can be taken independently and in any order.

- **Secure the extended workforce:** Give contractors and outsourced agents governed access on unmanaged devices, and stop shipping laptops.
- **Reduce the workplace envelope before 2027:** Move browser-first personas off managed devices and virtual desktops ahead of the workplace re-procurements, publish the few applications that genuinely need a remote session as applications rather than whole desktops, and cut the size of those contracts before they are let.
- **Retire point products:** Absorb data loss prevention, isolation, web gateway and privileged access into the browser, and reduce platform count.
- **Sequence the network decision:** Deploy the browser control point before scoping the 2027 network re-procurement, then commit to inspection capacity for the exceptions rather than for all user traffic.
- **Govern AI now:** Put guardrails around Copilot and public AI tools this quarter rather than after the first incident.

Buying it is straightforward. Island is available through G-Cloud as cloud software, which supports direct award. Under the Digital Assurance Playbook, a pilot sits comfortably inside delegated departmental authority. Nothing here needs a new mega-programme, a supplier relationship of the kind the Procurement Act was written to prevent, or a Treasury business case that outlives the government that commissioned it.

Start with one persona, one quarter, and a measured result.

Six questions worth asking before the Budget

- 1 How many of our users genuinely need a managed laptop, and how many need three web applications?
- 2 What does a virtual desktop seat cost us per year, fully loaded, and how many of those seats run nothing but a browser?
- 3 How long does it take to give a third-party agent secure access to citizen data, and how long to remove it?
- 4 How many separate products are we paying for that only ever act on a browser session?
- 5 If a member of staff pasted citizen data into a public AI tool this morning, would we know?
- 6 What share of the traffic our network inspection business case assumes it will read can it actually still read, and what will that share be in 2031 when the contract ends?

If the honest answers are uncomfortable, the useful part is that all six have the same fix, and none of them requires waiting for a migration to finish.

And five to press any vendor on, including us

- 1 Show me a reference customer of comparable scale and regulatory complexity, and let me speak to them without you in the room.
- 2 What is the availability record of your control plane, and what happens to policy enforcement when a client cannot reach it?
- 3 What is your patch cadence against upstream Chromium, and your turnaround commitment on a critical CVE?
- 4 Which UK certifications and assurance positions do you hold today, which are in progress, and where are the control plane and its data hosted for a UK customer?
- 5 What proportion of our traffic and our workflows will this genuinely not cover, and how would you propose we measure that before we commit?

A vendor that answers those five without hedging is worth a pilot. A vendor that does not is worth a different conversation.

The next twelve months

The reshuffle has handed departments more autonomy over their own technology decisions than they have had in a decade, and a Budget that will reward anyone who can show real savings quickly.

Departments that spend the next year waiting for the new machinery of government to settle will reach the next election with the same estate they have today.

The ones that change the control point will have banked the savings, secured the workforce they actually have rather than the one their device policy assumes, and made their eventual migration cheaper and safer.

**The applications can wait.
The access layer cannot.**

Customer references: TaskUs, yoummday, Teleperformance, Pfizer, Swiss Life, Landis+Gyr, Brightline Trains, Omada Health, Mattress Firm and University of the Pacific customer stories, published at island.io/customer-stories. Figures are published by Island, reported by the customers named, and not independently audited. US federal availability: Island press release, Billington Cybersecurity Summit, September 2024, and Carahsoft.

Sources: Legacy technology proportions and the £45bn productivity estimate: State of Digital Government Review, January 2025, together with the Commons Science, Innovation and Technology Committee response questioning the breadth of that figure. Central government IT maintenance spend: Government Security analysis. Supplier concentration, uncompleted awards and roughly £3bn in delays and resets across five major programmes: National Audit Office, Government's approach to technology suppliers, January 2025. Delegated departmental authority under the Digital Assurance Playbook, April 2026, and Procurement Act 2023 Competitive Flexible Procedure: gov.uk. Corporate by default device policy: security.gov.uk. Workplace and managed desktop contract expiry dates: Find a Tender and Contracts Finder. Network Services 3 (RM6116) scope and 17 July 2027 expiry: Government Commercial Agency. Encrypted Client Hello and TLS 1.3 inspection limits: IETF draft-camplng-ech-deployment-considerations, Microsoft Global Secure Access documentation, and AWS Network Firewall documentation. Cost modelling in this paper is illustrative, built on published public sector end-user compute benchmarks, and is not a price quotation.