National Federated Compute Services
NetworkPlus

The NFCS Roadmap

A Route to Federation for UK Digital Research Infrastructure

Version 1.2

The NFCS Lead Team

30th October 2026

UK Research and Innovation

The Lead Team

The NFCS Lead team consists of:

Prof. Jon Hays	Joint Project Lead
Dr Sadaf Alam	Joint Project Lead
Prof. Vassil Alexandrov	Project Co-Lead
Ms Josephine Beech-Brandt	Project Co-Lead
Prof. Adrian Hines	Project Co-Lead
Dr David Meredith	Project Lead Team Member
Mr Kieran Leach	Technical Coordinator
Prof. Emeritus Jeremy Yates	Technical Coordinator
Dr Rui Apóstolo	Community Engagement Officer
Mrs Gemma Hopton	Administration and Events Officer
Mr John Masih	Programme Manager

Acknowledgements

We would like to thank all those who have supported the NFCS NetworkPlus across the past two years, including all those who have taken part in our projects, attended workshops or conferences and otherwise joined in our community.

This work was funded through the UK Research and Innovation Digital Research Infrastructure Programme and the Engineering and Physical Sciences Research Council under Grant No. EP/Z534493/1.

AI disclosure: AI-based tools were used to assist with the editing and review of prose and the design of diagrams and presentation materials. The evidence, analysis, recommendations, and conclusions presented in this report derive from the NFCS consultation, the flexible fund projects, and the roadmap development process and are entirely those of the authors.

Contents

Contents ii

1 Executive summary 1

2 The Federation Vision 4
- 2.1 The case for federation . . . 4
- 2.2 The opportunity now . . . 4
- 2.3 Federation in practice . . . 7

3 From Evidence Base to Recommendations 9
- 3.1 From evidence to umbrella recommendations . . . 9
- 3.2 Prioritising the recommendations . . . 10
- 3.3 What the evidence could not answer . . . 11

4 The Umbrella Recommendations 15
- 4.1 Deliver first - the capabilities on which wider federation depends . . . 15
- 4.2 Enablers - early value alongside the foundations . . . 18
- 4.3 Deliver later - the next phase, once the foundations are in place . . . 20
- 4.4 Horizon capabilities - strategically important, longer term . . . 22

5 One federation, two roadmaps: shared foundations and coordination 24
- 5.1 Shared foundations and their dependencies . . . 24
- 5.2 Maintaining coherence across the roadmaps . . . 26
- 5.3 Delivery principles and assumptions . . . 28

6 Compute federation: Roadmap and pathways 30
- 6.1 Implementation options . . . 30
- 6.2 Appraisal . . . 32
- 6.3 The preferred pathway . . . 32
- 6.4 Delivery and verification . . . 35
- 6.5 Implications for future DRI funding activities . . . 36

7 Data services for compute federation: Roadmap and pathways 38
- 7.1 Implementation options . . . 38
- 7.2 Appraisal . . . 41
- 7.3 The preferred pathway . . . 42
- 7.4 Delivery and verification . . . 44
- 7.5 Implications for future DRI funding activities . . . 45

8 Roadmap to delivery and next steps 47
- 8.1 Propositions offered for challenge . . . 47

8.2 Activating the roadmap 49
8.3 Funding and sustaining federation 51
8.4 Final words 54

Bibliography **55**

Appendices **57**

Tracing Projects Across Recommendations **58**
Deliver First—see §4.1 58
Enablers—see §4.2 59
Deliver Later—see §4.3 59
Horizon—see §4.4 60

Flexible-funded Projects and their Outputs **61**

Areas of contention and compromise **63**

Glossary **64**



Chapter 1

Executive summary

The opportunity

The UK already operates a substantial digital research infrastructure, supported by major investments in compute, data, and AI. Researchers increasingly need to work across these services, yet often experience them as separate systems with different access arrangements, support models and approaches to data management. Federation offers a practical route to connecting this landscape without centralising ownership or replacing existing services.

It enables researchers to access resources more easily, providers to participate in a wider ecosystem without losing autonomy, and funders to realise greater value from national investment through coordination rather than duplication. By connecting existing capabilities through common standards, trust relationships, and operating frameworks, federation can increase the reach and impact of infrastructure that already exists while reducing barriers between researchers and resources.

Federation also strengthens the UK's sovereign capability to determine how its research infrastructure is organised and connected, reducing dependence on any single provider or technology choice while supporting full participation in the international infrastructure on which UK research relies.

> **"Federation works when everyone can use it."**
>
> — Surveying Accessibility of Fed. Compute for Junior Researchers (flexible-funded project).

The risk if we do not act

The UK is in the midst of a major new cycle of investment in DRI and AI capability, including the AI Research Resource, the rollout of National Compute Resources, and the Next National Supercomputing Service, together exceeding £2 billion, that presents a unique opportunity to embed federation from the outset. Establishing common approaches while services and operating models are still evolving is substantially easier than attempting to retrofit them later. Delaying action risks embedding barriers that become progressively harder and more expensive to remove.

What the evidence says

Across the evidence base, a consistent picture emerged. A small number of capabilities underpin wider federation: federated identity and access management , coordinated cybersecurity, a federation governance and operating model, a common data-movement, access and storage standard, a single front-door for onboarding and support, and sustained attention to culture, participation and user needs.

The NFCS programme funded 26 projects between July 2025 and July 2026, spanning community, governance and technology. They generated 259 recommendations consolidated into six themes and 17 umbrella recommendations, with full traceability to the original evidence, and prioritised on both urgency and strategic importance. These are organised into four bands: those that must begin first, those that can progress alongside them, those that depend on the foundations being in place, and longer-term horizon opportunities.

The programme also identified a small number of unresolved questions. The most significant of these concern: future storage demand and costs, long-term responsibility for data retention, sustainable operating models for shared federation services, what it should mean in practice for participants to be part of federation, and what agent-mediated use of research infrastructure requires of it.

One federation, two roadmaps

UKRI asked for separate roadmaps for compute and data-services. These are not independent programmes, but complementary pathways within a single programme. Some capabilities need to be developed once and used by both: common approaches to identity and access, governance and allocation, data movement and discovery, and secure environments for sensitive data.

The two roadmaps depend on the same foundations that need to be developed once and used by both. Compute federation begins with identity, governance, cybersecurity and the interoperability needed for services to participate with confidence. Data services begin with common standards that allow data to be discovered, accessed, moved and secured across independently governed providers, while future demand and capacity are measured alongside delivery. Together these roadmaps provide a practical route to federation that builds on existing investments, reduces fragmentation, and allows adoption to grow incrementally.

What we recommend

Each roadmap appraises five implementation options, including a do-nothing baseline. The evidence favours a *foundation-led* pathway for compute federation and a *standards-led* pathway for data services. Both depend on shared approaches and the preferred pathways are presented as evidence-based recommendations for discussion with research communities, providers, and funders, and as the basis for the development of a delivery plan.

Delivery should be phased and evidence-led. The first phase focuses on both the foundations of federation and key shared approaches: identity, governance, cybersecurity, data interoperability, onboarding, and support. Later capabilities are introduced only when those foundations have been shown to work in practice and adopted by participating providers, with investment released against demonstrated adoption in production. The first user-visible benefits arrive within that first phase. Researchers gain simpler access to participating services, clearer routes to support,



and improved ability to discover and move data between providers. Providers gain common standards, tested implementation patterns, and a clearer route into federation.

Investment and Sustainability

Federation is not principally about creating new infrastructure. Its purpose is to increase the value realised from existing infrastructure and future investments. By improving interoperability, reducing duplicated effort, and broadening participation, federation allows national and institutional investments to function more effectively as a connected system.

The evidence therefore supports funding federation foundations as infrastructure instead of short-lived projects. Identity, governance, cybersecurity, testing environments, and coordination all carry continuing obligations that persist after capability development is complete. Funding models should reflect that reality. Combining the preferred compute (80 FTE-years) and data-services (48 FTE-years) pathways gives an indicative activation effort of around 128 FTE-years across the three phases, approximately £16 million of effort, against the £2 billion already committed to the infrastructure it would connect. Activation means establishing the coordination function, working groups, and testing environments needed to develop the shared foundations, while progressing the preferred compute and data pathways through phased delivery. Several approaches to funding and governance are possible. The roadmap does not prescribe a single organisational model, but shows how the recommended approach can be delivered through proven coordination mechanisms already used in the UK and internationally.

> **"Federation succeeds when project context, authority and lifecycle operations are funded as part of the infrastructure — not bolted on afterwards."**
>
> — Bridging the Gap: Aligning Project Admin with Access to DRI (flexible-funded project).

Conclusion

The evidence suggests that federation is not primarily a technology problem. The UK already possesses many of the technical, governance, and community foundations required. The challenge is to coordinate, sustain, and grow those foundations into a coherent national capability. Federation succeeds when technical interoperability is supported by enduring governance, coordination, and community engagement. This brings together providers, funders, and research communities to maintain the shared foundations, and support adoption.

This roadmap proposes an evidence-based, community-driven route for doing so, while making explicit the assumptions, uncertainties, and decisions that remain open for challenge. The immediate decision is whether to activate that route, and on what terms.





Chapter 2

The Federation Vision

Our vision is a UK digital research infrastructure that operates as a system of systems: researchers can find and use the compute, data and AI capabilities they need across institutional boundaries; providers can participate in a wider national offer while retaining control of their infrastructure; and funders can coordinate investment across the landscape while effectively targeting areas of unmet need.

2.1 The case for federation

> **"The federation of geographically distributed computing resources offers radical enhancements in accessibility, resilience, and collaborative capability."**
>
> — Federated Edge-HPC Architectures for AI Workflows in Privacy-Sensitive and Real-Time Domains (flexible-funded project).

The UK already operates a substantial digital research infrastructure, built through investments by universities, national facilities, research councils, and government. The diversity of this infrastructure is a strength, but as many of these systems have largely developed independently, researchers often encounter them as separate systems with different identity and access arrangements, policies, support models and approaches to data management. Researchers can struggle to discover appropriate services, move workloads and data between different services, or combine access to different resources to support more complex workflows.

Federation provides a practical response to this challenge. It does not require centralised ownership or the replacement of existing infrastructures. Instead it provides the shared standards, trust relationships and governance needed for services to work together.

For researchers, this means simpler access to a broader range of capabilities and fewer barriers to access. For providers, it means greater visibility and participation without loss of autonomy. For funders, it provides the means to coordinate investment, identify gaps in provision, and increase the value derived from existing infrastructure.

Some of the key challenges identified across the NFCS evidence base, and the ways in which federation addresses them, are summarised in Table 2.1.

2.2 The opportunity now

The case for federation is not new. The need for greater coordination across the UK's digital research infrastructure landscape has been identified repeatedly in national reviews and strategies. The Government Office for Science review of large-scale com-

Challenge	Federation response
Researchers require different identities, access arrangements and support processes across services.	Federated identity and access management, common on-boarding approaches, and shared support pathways
Data is difficult to discover, access, and move between providers	Common standards for data movement, cataloguing, access, and storage
Services are difficult to discover, particularly for new users and emerging communities.	A single front-door, improved service discovery, and coordinated user support
Providers develop similar capabilities independently, increasing duplication of effort	Shared standards, reference implementations and coordination mechanisms to allow reuse of solutions.
Trust, accountability, and security approaches vary or are incompatible between organisations	A shared governance and operating model together with coordinated cybersecurity approaches
Funders have limited visibility of gaps, duplication, and opportunities for coordination	Federation-level governance, allocation, and reporting mechanisms that support strategic investment decisions.
Research communities differ in their readiness, skills, and awareness of digital research infrastructure	Communities of practice, skills development and user-needs driven federation activities.

Table 2.1: The challenges federation seeks to address

puting (1), the *Independent Review of the Future of Compute* (2), UKRI's research and innovation infrastructure roadmap (3), the DSIT *UK Compute Roadmap* (4), and JISC's *Mapping Federation Journeys for Optimising the UK Digital Research Infrastructure* (5) have highlighted the costs of fragmentation and the need for closer coordination across services and investments. This roadmap builds on this body of work to provide an evidence base and delivery pathways for how federation could be realised in practice.

The timing is now particularly significant. Large investments are being made across the UK's digital research infrastructure, including the AI Research Resource (AIRR) (6), National Compute Resources (NCR) (7) and the Next National Supercomputing Service (NNSC) (8). Without effective coordination there is a danger that new capabilities are operated as isolated assets, increasing complexity for users and reducing the overall value of the investments.



> **"Federated data and compute services have an increasingly important role in supporting interdisciplinary environmental research. In practice, users often experience these services as fragmented, difficult to navigate and disconnected from their needs."**
> — UNITED (User Needs Informed Technologies for Environmental Data) (flexible-funded project)

At the same time, researchers increasingly require digital research infrastructure to function as a connected service, not a collection of independent systems. Modern research projects often combine data, compute, and AI capabilities provided by different organisations and infrastructures, while also needing to integrate with local institutional systems. Without effective coordination, researchers bear the cost of navigating the resulting complexity.

Federation is also strategically important. Decisions taken in the coming investment cycle will influence how the UK's digital research infrastructure develops over the next decade. Federation strengthens the UK's sovereign ability to shape how its infrastructure is organised and how it evolves. Without it, those outcomes emerge implicitly from individual investments and technology choices. Similar concerns are increasingly visible internationally. Canada's Digital Research Alliance explicitly links coordination of digital research infrastructure to national data sovereignty (9), while the European Union frames comparable challenges through the concept of strategic autonomy: the ability to cooperate internationally while retaining sufficient control over critical technologies and infrastructure (10).

International experience also demonstrates that federation can be successful without centralising ownership or control of infrastructure. Established initiatives including EuroHPC (11), EOSC (12), ELIXIR (13), and the Worldwide LHC Computing Grid (14) demonstrate that services can be connected through shared standards and governance while retaining local autonomy. More recent programmes, including the US National Science Foundation's ACCESS (15) programme and the Department of Energy's Integrated Research Infrastructure (16) programme, reflect a growing recognition that national capability increasingly depends on connecting existing services rather than creating new silos. Germany's National Research Data Infrastructure (nfdi) (17) demonstrates a complementary approach in which discipline-led consortia, and coordinated governance are used to develop sustainable national research-data capabilities across a diverse research landscape.

Importantly, federation need not be built from scratch. The UK already possesses many of the technical, organisational, and community foundations required. These include established identity federations such as the UK Access Management Federation and its inter-federation with eduGAIN (18), community-operated infrastructures, collaborative research programmes, and national coordination initiatives such as IRIS (19), which has demonstrated the value of bringing research communities, providers, and funders together around shared infrastructure challenges. Similarly, DiRAC (20) demonstrates how previously independent competing research commu-



nities can converge around common governance and shared infrastructure to deliver a broad science programme, while retaining strong institutional identity. The UK also participates in internationally connected research infrastructures such as GridPP's participation in the Worldwide LHC Computing Grid (14, 21). Together, these demonstrate that many of the capabilities, relationships, and governance models already exist in some form. The challenge is not to replace them, but to learn from these examples to connect things more effectively.

The opportunity is therefore immediate. Many of the building blocks already exist, but integrating them while infrastructure and services are still evolving is substantially easier than attempting to retrofit later. Delaying coordination risks embedding technical and organisational barriers that become progressively harder and more expensive to remove.

2.3 Federation in practice

Federation is not a single intervention or investment. It does not require every aspect of infrastructure to be standardised or centrally coordinated. Its purpose is to focus on the areas where common approaches create the greatest benefit, while preserving diversity where it provides value.

The NFCS Programme approaches federation through three complementary pillars: Technology, Governance, and Community. The technology pillar considers the technical mechanisms that allow services to interoperate. The governance pillar addresses trust, accountability, security, allocation, and the operating processes and policies needed to sustain federation over time. The community pillar focuses on participation, on-boarding, skills, culture, and ensuring that federation evolves in response to the needs of research communities.

Federation is not an all-or-nothing choice but a progressive model of participation. At one end of the spectrum, services may remain independent with no shared standards or processes. At the other, they may expose resources, data and operations directly through federation mechanisms. Between these extremes, many services may choose to adopt common standards for identity, access, and interoperability without substantial changes to how they are operated.

Federation participation can be described through a set of levels. These represent increasing degrees of participation in federation rather than maturity levels or stages that every provider must progress through. Organisations may participate at different levels, with different services according to need, maturity and operational constraints:

- Level 0 (Independent): services operate entirely using local arrangements
- Level 1 (Interoperable): services adopt common standards for identity, access and data, making them discoverable and usable within the federation.
- Level 2 (Integrated): services expose resources, data and operational capabilities directly through federation mechanisms to provide integrated services.



Future participation models may also need to accommodate delegated and agent-mediated interactions.

The practical question is therefore not whether to federate, but what should be federated and to what degree. Chapter 5 builds on this model by identifying the shared foundations and coordination mechanisms needed to support participation at each level.

> **“Try to reach audiences you don’t know much about. They may surprise you.”**
> — Federating Compute at Scale: Power/Knowledge Narratives and Imaginaries of the National Federation of DRIs (Workshop insight)



Chapter 3



From Evidence Base to Recommendations

> **"The findings suggest that decisions about federated infrastructure design are decisions about participation."**
>
> — Inclusive Futures: User stories and Mapping Pathways for National Federated Compute Services (A flexible-funded project)

Between July 2025 and July 2026, the NFCS programme funded 26 projects spanning community, governance, and technology aspects of federation. The outputs of these projects, together with evidence gathered through workshops, consultations and community activities provided the foundation for the compute and data federation roadmaps presented in this document. This chapter explains how that evidence was synthesised into 17 Umbrella recommendations, prioritised and used to identify both the foundations of federation and the knowledge gaps that remain.

3.1 From evidence to umbrella recommendations

The NFCS Programme has generated evidence from across an unusually broad range of stakeholder groups. Together they represent both highly federation-mature research communities and emerging research communities that are not yet fully engaged with national research computing.



Figure 3.1: From project evidence to prioritised recommendations.

Collectively the flexible-funded projects generated 259 recommendations. This breadth provides confidence that the roadmap reflects a wide cross-section of stakeholder perspectives. However, it creates a practical challenge that no funding programme can act effectively on hundreds of individual recommendations.

The evidence was organised and synthesised through four successive layers:

Pillars give breadth, **themes** give structure, **umbrella recommendations** provide consolidation, and **prioritisation** provides sequence.

First each recommendation was tagged against one or more of the programme's three pillars: Community, Governance, and Technology, that provide a check on the breadth that no dimension of federation was overlooked.

Second, recommendations were grouped into six themes (A-F), outlined in Table 3.1. These provide a navigable structure for the evidence and reflect the reality that many projects span multiple pillars. They identify the coherent pathways through which federation can be developed. Third, the 259 recommendations were consolidated into

Theme	Covers
A. Access & Orchestration	Reaching and moving workloads across federated compute — workflows, job orchestration, tokenised access.
B. Data Movement & Access	Moving, sharing, storing and curating data across sites — common APIs, encrypted/TRE storage, and data contracts.
C. Software, AI & Emerging Platforms	Software and AI container supply chain, governed AI research services, quantum–HPC integration.
D. Governance & Cybersecurity	Identity and access, cybersecurity, operating model.
E. Community, Skills & Support	Onboarding, admin, support and the single front door, skills, and RSE capacity, culture, and user needs.
F. Sustainability & Fair Allocation	Fair, transparent resource allocation and accounting; carbon- and resource-aware federation.

Table 3.1: The six themes (A–F) that structure the recommendations.

17 umbrella recommendations. This allowed related recommendations to be brought together as coherent interventions while retaining full traceability to the underlying project evidence.

Finally, the umbrella recommendations were prioritised using a priority / urgency framework and placed into a 2×2 matrix. This is then used to build a structured set of recommendations into a sequenced roadmap, identifying the foundations that must be delivered first and the capabilities that can follow.

The themes, umbrella recommendations, and prioritisation model were reviewed through the NFCS governance and communities structures, including the NFCS Community board, the flexible-funded projects at the July 2026 project recommendations and review meeting, and the wider community at the September 2026 conference on the draft roadmap. Feedback from these groups directly informed refinement of the roadmap and provides confidence that the resulting framework reflects both the evidence gathered and the perspectives of the communities that will ultimately deliver and use federation. Key areas of contention or compromise arising from that consultation, even where these sit outside this main report, are recorded in Appendix 8.4.

3.2 Prioritising the recommendations

The prioritisation identifies a relatively small number of foundational capabilities that are both urgent and strategically important. These form the deliver first band. Other recommendations either provide early value alongside these foundations, depend upon them, or represent longer-term strategic opportunities. This prioritisation provides the structure for the recommendations chapter that follows. Diagrammatically,



Ref	Umbrella recommendation	Theme	Priority	Urg.	Quad.
A1	Federated Workflow & Job Orchestration	A	High	Later	②
A2	Federated Access Mechanisms	A	Moderate	Now	③
A3	Industry & Public Sector Access	A	High	Later	②
B1	Common Data Movement, Cataloguing, Access & Storage Standard	B	High	Now	①
B3	Federated Trusted Research Environments	B	Moderate	Now	③
C1	Federated Software & Container Supply Chain	C	High	Later	②
C2a	AI Research-Support Services across the Federation	C	High	Later	②
C2b	Responsible AI governance	C	High	Later	②
C3	Quantum–HPC Integration Pathway	C	Moderate	Later	④
D1	Federated Identity & Access Management	D	Highest	Now	①
D2	Coordinated Cybersecurity Governance & Operations	D	Highest	Now	①
D3	Federation Governance & Operating Model	D	High	Now	①
E1	Single Front Door, Onboarding, Admin & Support	E	Highest	Now	①
E2	Skills, RSE & Communities of Practice	E	Moderate	Now	③
E3	Culture, Widening Participation & User Needs	E	High	Now	①
F1	Fair, Transparent Resource Allocation & Accounting	F	High	Later	②
F2	Carbon & Resource Aware Federation	F	Moderate	Later	④

Table 3.2: The 17 umbrella recommendations. **B2**, covering storage, cataloguing and curation, was merged into **B1** following the July 2026 consultation, on the grounds that separating movement from storage and cataloguing made an artificial boundary; identifiers were not renumbered, so traceability to the underlying evidence is preserved. **C2** was split into **C2a** and **C2b**.

the prioritisation is shown in Figure 3.2. The matrix identifies six recommendations that must begin first, three that can progress alongside them, six that follow once the foundations exist, and two longer-term opportunities.

3.3 What the evidence could not answer

The programme's evidence is broad, but it is not complete. Some questions could not be settled from surveys, workshops and short technical projects, either because they require operational data that is not currently collected, because they depend on decisions that have not yet been taken, because they fall outside the scope the programme was funded to address, or because they emerged after the evidence base was commissioned. These are recorded below rather than left implicit. Each is given an identifier so that it can be cited from the roadmap chapters and, where appropriate, lifted into a future funding activity.

Several of these gaps are closed by the mechanisms described in Chapter 5: the testing environments allow questions of implementation to be settled empirically, and the



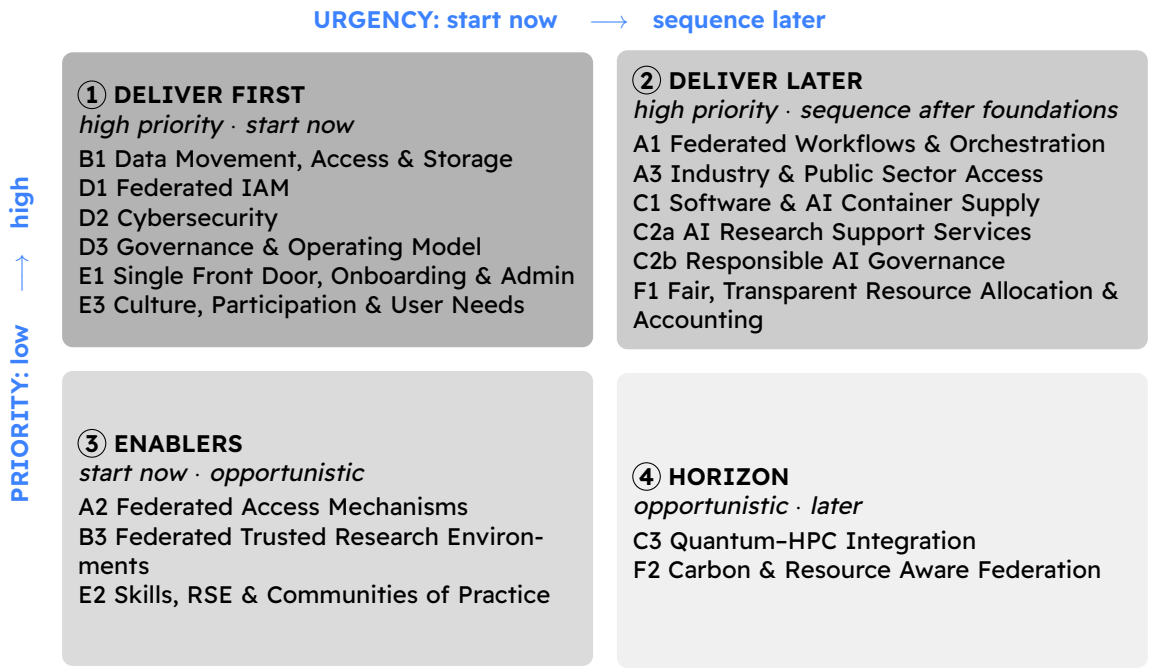


Figure 3.2: The 17 umbrella recommendations on the priority × urgency map. Governance (D1–D3) sits early in quadrant ① as an enabling band but does not gate everything.

working groups are the route by which questions of definition and operating model are resolved. Others require work that no part of the current programme can undertake, and are the clearest candidates for future DRI funding activity.

Table 3.3: Knowledge gaps carried forward from the NFCS evidence base. **fed-g7** arises from feedback on the draft roadmap at the September 2026 conference.

ID	What the evidence does not establish	Why it matters	What would close it
Compute federation			
fed-g1	What "federated" requires of an existing service in operational terms.	The term is widely used, including in national strategy, without a definition against which a service can be assessed. Without one, conformance cannot be verified.	A definition and maturity model agreed through the federation working groups and tested against participating services.
fed-g2	Which identity and authority models institutions are able to implement in practice.	Requirements can be specified on paper, but institutional identity provision varies widely and adoption is voluntary.	The identity integration testbed, working with early-adopter sites across more than one council domain.

continued overleaf

Knowledge gaps, continued

ID	What the evidence does not establish	Why it matters	What would close it
fed-g3	What conformance to the data-movement and access standard means in practice.	Conformance is the entry profile for federation; it cannot be required of providers until it is defined and demonstrably testable.	A conformance suite developed and operated in the testing environments.
fed-g4	How portable workloads are across heterogeneous schedulers and site configurations.	Workflow portability is a principal user-visible benefit of federation and is assumed by several later capabilities.	Phase 2 orchestration work, evaluated across sites spanning DRI types.
fed-g5	The operating model and cost-recovery arrangements for the shared foundations.	The foundations carry recurrent obligations. How they are governed, and who pays for them over time, is unresolved.	Governance and accounting working group, informed by existing federation practice in the UK and internationally.
fed-g6	What federation readiness means for a domain or community starting from close to zero.	Research councils and communities differ markedly in federation maturity; provision designed around mature communities would exclude the rest.	Targeted requirements work with less-established communities, building on the programme's engagement projects.
fed-g7	What agent-mediated use requires of federation.	Identity, allocation, accounting and the safe-people model assume an identifiable human actor. Agent-mediated use is already occurring on production services, under user credentials, outside these frameworks.	The identity and governance working groups, with delegated-authority and interface profiles exercised in the testing environments.
Data services			
sto-g1	The storage needs of UK computational research beyond those the federation itself requires.	This is the wider question posed alongside federation. The programme addressed the data services federation depends on; the general case fell outside its scope.	A dedicated assessment, commissioned separately from the federation roadmap.
sto-g2	Data volumes, growth rates and unit costs across the UK research estate.	No capacity decision — national, institutional or commercial — can be sized without them. Every option in the data roadmap is weakened by their absence.	The commissioned demand study, funded as a track alongside the standard rather than in place of it.

continued overleaf

Knowledge gaps, continued

ID	What the evidence does not establish	Why it matters	What would close it
sto-g3	Who funds retention and preservation once a grant ends.	Custody obligations attach to data at creation and outlive the project that produced it. There is at present no answer.	A funder-level decision, informed by the governance and accounting working group.
sto-g4	How far metadata and cataloguing practice can be made to conform across disciplines.	Without it, federated data is movable but not findable. Discipline practice diverges more here than anywhere else in the data roadmap.	The metadata and cataloguing conformance strand, tested across discipline repositories.
sto-g5	Egress and tiering economics across a mixed institutional, national and cloud estate.	These costs determine whether federated data movement is affordable in practice, and they cannot be settled on paper.	The data-transfer testbed, operated with users and providers.



Chapter 4

The Umbrella Recommendations

The process used to synthesise the evidence base into a manageable set of recommendations was described in Chapter 3. This chapter discusses each of the umbrella recommendations grouped by the four recommendations bands derived from the priority-urgency framework. Together these recommendations provide the foundation for the roadmap pathways, sequencing decisions, and delivery options in subsequent chapters.

4.1 Deliver first - the capabilities on which wider federation depends

The evidence consistently identified a small set of capabilities that are both urgent and strategically important. Together they form the deliver-first band: the standards, trust relationships, governance arrangements, and user pathways that allow independently operated services to function coherently. These capabilities should begin early, although progress elsewhere need not wait for every element to be completed. Some also form part of the shared foundations of the compute and data-services roadmaps as discussed in Chapter 5.

B1: Common Data Movement, Cataloguing, Access, and Storage Standards

> **"Federated Data Movement is important because it removes access barriers from users, enabling them to move data and control and schedule data movement without needing specific permissions or bespoke support from system operators."**
>
> — Federated Data Movement (A flexible-funded project)

> **Recommendation:** Establish a standard set of data-transfer mechanisms operating at different scales, usable without manual intervention and secure for national and international access, with roles and policy-based access controls for different data classifications; service catalogues, potentially AI-assisted for interoperability, evolving into a searchable federation data catalogue integrated with Federated IAM for both access and citation tracing; and a national strategy for data contracting aligned with the UKRI Data Policy, covering open, sensitive and commercial data, with standardised templates and a possible national data advisory service.

Data underpins research computing, making the ability to discover, store, access, and move data fundamental to federation. Technical interoperability alone is insufficient. Differences in access policy, data classification, contracting, and institutional responsibility can obstruct access or movement even where compatible transfer mechanisms exist. B1 brings the technical, policy, and governance dimensions of data interoperability into a single umbrella recommendation.

D1: Federated Identity and Access Management

> **"A consistent and joined up IAM system is important where users directly interact with many different systems otherwise users can be overwhelmed by different procedures and to handle different authentication credentials."**
>
> — Federated IAM for Existing Infrastructures (A flexible-funded project)

Recommendation: Establish a harmonised Identity and Access Management framework for UK DRI, compliant with the AARC Blueprint Architecture (22) for interoperability and capable of evolving to meet new technologies and diverse research communities' requirements, covering persistent identifiers across a career, virtual organisations, inter-federation bridging, and an identity provider of last resort, underpinned by a governance regime with paid support contracts, SLAs, resilience and operational targets.

A key repeated theme across projects has been the necessity of a **Federated Identity and Access Management Framework**. This is both a direct user benefit and an enabler of wider federation. It reduces the need for researchers to navigate separate credentials and access procedures, while providing services with trusted information about identity, authorisation, and project context. Its value therefore lies not just in enabling single sign-on, but also supporting consistent and accountable access across participating services and underpinning several other recommendations in the roadmaps.

D2: Coordinated Cybersecurity Governance and Operations

> **"Research computing has unique characteristics which make it difficult to impose historically enterprise focussed Cyber Security approaches. Whilst DRI Federation reduces the friction of using the DRI for researchers, it also opens up potential vulnerabilities which need addressing as part of the overall DRI Federation roadmap."**
>
> — DRI Cybersecurity (A flexible-funded project)

Recommendation: Develop a coordinated cybersecurity governance and risk-management approach, interoperable with the federation governance model, on a shared baseline aligned to the NCSC Cyber Assessment Framework (CAF) (23), proportionate for organisations without full compliance capacity, supported by training for users and providers, with federation-level incident tracking and real-time information sharing through a federation-level Cyber Security Incident Response Team (CSIRT).

Federation increases the need for coordinated cybersecurity because vulnerabilities, incidents, and assurance requirements can extend across organisational boundaries.



A shared baseline would provide confidence to users and organisations while allowing implementation to remain proportionate to organisational capacity. Federation-level coordination would support consistent risk-management, incident reporting, and information sharing without requiring all participating organisations to operate identical security arrangements.

D3: Federation Governance and Operating Model

> **"Inclusive federated infrastructure futures are shaped more by governance design than by technology choices alone."**
>
> — Inclusive Futures: User stories and Mapping Pathways for National Federated Compute Services (A flexible-funded project)

Recommendation: Establish a federation governance and operating model that owns the strategic purpose of federation, mandates shared tools and frameworks, monitors risk and fosters trust, and is accompanied by a federated culture and learning observatory, with representation reviewed as communities change.

A consistent theme emerging from the evidence is the need for federation governance. Technical standards alone are not sufficient to sustain federation. A governance and operating model is needed to establish decision-making authority, define responsibilities, resolve disagreements, maintain shared frameworks, and provide accountability across organisational boundaries.

It provides ownership of the rules, definitions, and processes that underpin federation, including what it means to participate and how conformance is assessed. In doing so, it establishes the institutional basis through which the technology, community, and governance aspects of federation can evolve and remain coherent over time.

E1: Single Front Door, Onboarding, Admin and Support

> **"The project found that the principal barrier is not the absence of identity, workflow or access-control tools. It is the weak connection between project authority, institutional processes and the full access lifecycle."**
>
> — Bridging the Gap: Aligning Project Administration with Access to Digital Research Infrastructure (A flexible-funded project)

Recommendation: Establish a user-centred single front door, potentially leveraging AI techniques: a guided route, with plain-language and accessible explanations, to what services exist and what they cost, a searchable UK DRI platform, clear support routes, a minimum national project-to-access information profile, reusable policy defaults that distinguish operational delegation from institutional accountability with auditable models, lifecycle communications, and cross-council communities of practice.



Onboarding, administration, and support are part of the infrastructure researchers experience. Differences in service discovery, project approval, authorisation, access procedures, and lifecycle management create barriers and delays for users, and repeated administrative effort for providers and funders. The single front door must therefore connect these activities through common processes and information flows, rather than operate solely as a portal layered over otherwise fragmented arrangements.

E3: Culture, Widening Participation and User Needs

> **"High-Performance Computing (HPC) and federated digital research infrastructure are increasingly central to data-intensive research, simulation, artificial intelligence, digital twins, and evidence-led decision-making. Yet the ability to benefit from HPC remains uneven across disciplines, sectors and institutions"**
>
> — Enhancing HPC Adoption Through User-Centred Design (A flexible-funded project)

> **"Our aim is to engage with the A&H communities to collectively elicit requirements around federation within the current DRI landscape to support the NFCS roadmap."**
>
> — Federation in the Arts & Humanities (A flexible-funded project)

> **Recommendation:** Map recurrently how different research communities understand federation; take a community-led approach to widening participation; accommodate non-linear and iterative ways of working; provide capabilities beyond traditional research computing; and align community gateways with the front door.

The evidence base included research communities with widely differing levels of experience, readiness, and engagement with national research computing. Their needs cannot be addressed through technical interoperability alone. Federation must respond to differences in research culture, institutional capacity, skills, and expectations. This requires mechanisms and routes through which a broad range of research communities, particularly those that have not traditionally engaged with national research computing, can continue to shape the developing infrastructure.

4.2 Enablers - early value alongside the foundations

The enabler recommendations can progress alongside the deliver-first activities. They offer early value, help test aspects of federation in practice, and reduce risk for capabilities developed later. Although they themselves are not part of the deliver-first band, they support adoption and contribute to the shared capabilities on which the wider federation will depend.

A2: Federated access mechanisms

> **Recommendation:** Build upon established access technologies and production federation experience to develop interoperable approaches for authorisation and access across federated services. Existing approaches, in-



cluding token-based authentication and authorisation mechanisms should be leveraged wherever possible. Develop and evaluate reference implementations spanning multiple DRI environments, ensuring that different access approaches can be harmonised through common policy, identity, and governance frameworks.

Federated identity establishes who a user is and what authority they hold. Access mechanisms provide the means through which that authority can be exercised across participating services without requiring service-specific credentials or bespoke local account management. Proven approaches exist within large-scale federated heterogeneous infrastructures, such as the Worldwide LHC Computing Grid. Federation should build on this and similar experience while recognising that different services may adopt different access technologies. The goal is therefore not a single mandatory implementation, but interoperable approaches that allow authority, permissions, and policy to be applied consistently across federated compute, data and storage services.

B3: Federated Trusted Research Environments

Recommendation: Build upon existing work on secure federated access to sensitive data, including approaches emerging from DARE UK and related TRE initiatives, to develop and evaluate a common approach to interoperable encrypted storage patterns for federated Trusted Research Environments. Evaluate and provide reference implementations that support portability and interoperability between TREs, while retaining local control of cryptographic keys, governance, and compliance obligations.

Trusted Research Environments are an increasingly important component of UK digital research infrastructure, providing controlled access to sensitive and regulated data. Significant work is already underway through DARE UK and related activities to understand how secure research environments can interoperate whilst maintaining appropriate governance and assurance. Federation creates a need for compatible security patterns that allow data and workloads to move between participating environments without weakening local control. The goal is therefore not a single mandatory storage platform or security implementation, but interoperable approaches that provide a trusted foundation for sensitive data services operating across a federated infrastructure.

E2: Skills, RSE and Communities of Practice

Recommendation: Establish a coordinated national skills and community-development programme that builds upon existing communities of practice, professional networks, and training initiatives to support the adoption, operation, and use of federated services.



Federation depends on people as much as technology. Researchers, facilitators, and technical professionals need the knowledge and skills required to use, support and operate federated services. The evidence identified a strong existing ecosystem of training activities, professional networks and communities of practice. Federation should build on these existing activities, adding effort only where interoperability, governance and operational practice introduce new requirements.

4.3 Deliver later - the next phase, once the foundations are in place

The deliver-later recommendations represent strategically important capabilities whose full implementation depends on progress in the delivery-first band. Their placement doesn't imply that all preparatory work should be postponed. Requirements gathering, standards development, and technical exploration can begin earlier, but deployment at scale should be sequenced after the relevant foundations are sufficiently mature.

A1: Federated Workflows and Orchestration

> **Recommendation:** Review the workflow and job-orchestration technologies used by research communities across DRI, demonstrate and document their integration with federation, use workflow platforms to provide community sandboxes for portable reproducible pipelines, and integrate workflow platforms with data producers such as laboratories.

Research workflows increasingly span multiple datasets, software environments, organisations, and compute services. Differences in schedules, access arrangements, data services, and local configurations can make those workflows difficult to deploy across providers. Federation should support interoperable workflow and orchestration technologies that allow researchers to combine services from multiple providers more effectively. Delivery depends on progress in identity and access management, federated access mechanisms, governance and data interoperability, making this a natural follow-on capability once the federation foundations are established.

A3: Industry and Public Sector Access

> **Recommendation:** Conduct requirements gathering for commercial, industrial, and priority public-sector users; provide guidance and governance for sensitive commercial data; document the shared cybersecurity baseline so that the level of assurance is clear; support these users at the single front door; and support interoperability with cloud providers.

Federation has the potential to extend the benefits of the UK DRI beyond publicly supported research, supporting collaboration with industry and the public sector. These users often have requirements that differ from those of academic researchers, including commercial confidentiality, assurance, contracting, and procurement, as well as interoperability and support expectations. These should not be assumed to



align automatically with existing academic access models. Targeted requirements gathering is needed before access arrangements are scaled, ensuring wider participation can be supported without weakening security and accountability, or the needs of existing research communities.

C1: Federated Software and Container Supply Chain

> **Recommendation:** Map the supply-chain characteristics of existing DRI capabilities, including ownership, licensing and nation-state involvement; apply a review process to new components; deploy a federated container registry with cache registries at constituent DRIs; and set policies for container lifecycle and unmaintained images.

Federated infrastructure increasingly depends on software, containers, and services developed and maintained by a wide range of organisations. In many environments, software deployment and lifecycle management are becoming increasingly delegated to users and projects rather than being solely managed by infrastructure providers. This creates challenges around software assurance, provenance, licensing, maintenance, and supply-chain security, particularly where software may be deployed across multiple independent providers. Shared review processes, federated registries, and agreed policies can reduce duplicated effort while increasing confidence in software used throughout the federation. This complements the wider cybersecurity recommendation, which addresses risks in infrastructure operations, by covering those originating in software supply chains and delegated deployment models.

C2a: AI Research Support Services

> **Recommendation:** Provide an on-premise or in-federation LLM facility with configurable knowledge-base sources, back-end and access-control policies; support DRI sites in deploying AI helpdesk facilities; and support users in large-scale data analysis and researchers with agentic coding assistants.

Federated digital research infrastructure has the potential to be difficult to navigate, especially for new users and research communities with limited previous engagement with national DRI. AI-supported services have the potential to improve on-boarding, support, and engagement by helping users discover relevant services, search and interpret documentation, understand polices and procedures, and access guidance in forms appropriate to their experience and needs. These capabilities may help maintain effective support that can scale as federation grows and human support resources become more distributed. Federation provides a framework through which shared knowledge bases and AI-enabled support services can be developed collaboratively.



C2b: Responsible AI Governance

> **Recommendation:** Develop a federation-wide governance framework for AI integration, adoption and user support, with a dedicated security work stream tracking agentic AI and other emerging technologies, safeguards against large-scale data scraping, and a federation-wide documentation standard suitable for AI consumption.

Services deployed across a federation raise governance challenges that extend beyond any individual provider. Differences in approaches, policy, and operations can undermine trust and interoperability if managed inconsistently, particularly where sensitive information is involved. A federation-wide governance framework would establish common expectations while allowing providers flexibility in their choice of technologies and implementation approaches. Agentic AI introduces additional challenges because actions may be delegated to autonomous systems operating across multiple services, providers, and data sources, requiring clear authority, accountability, and auditability. Consultation indicated that such use is already occurring ahead of the frameworks intended to govern it; the requirements this places on federation are recorded as **fed-g7**.

F1: Fair and Transparent Resource Allocation and Accounting

> **Recommendation:** Establish a unified data collection standard for DRI capturing a minimal agreed schema; define a set of standard metrics reported alongside occupancy-based accounting; and establish a cross-system benchmarking and exchange-rate framework so that allocations are portable across heterogeneous systems.

Federation requires allocations, usage, and capability to be consistently understood across heterogeneous services. Resources provided by different services may not be directly equivalent, making it difficult to compare, transfer, or account for allocations. A common set of metrics, reporting standards and benchmarking are needed to support transparency, and portability of allocations across the federation to provide the basis for fair and evidence-based resource management.

4.4 Horizon capabilities - strategically important, longer term

The horizon recommendations represent strategically important areas that may shape the longer-term evolution of a federated infrastructure. Their successful progression into wider delivery depends on the maturity of the federation foundations, viable operating models, and demonstrated community need. They should remain under active review. They are neither immediate priorities nor candidates for indefinite deferral.



C3: Quantum-HPC Integration

> **Recommendation:** Define and demonstrate a minimum viable federated quantum-DRI access service alongside reference hybrid quantum-HPC and quantum-AI architectures; provide simulator-first entry pathways with common programming and middleware expectations; and incorporate quantum literacy into federated skills, training and governance.

Quantum technologies are likely to be introduced as specialised capabilities operating alongside conventional compute, AI, and data services. Federation provides a framework through which these capabilities can be integrated into the broader DRI ecosystem. Early preparation can reduce future integration costs while enabling research communities to experiment with quantum technologies before large-scale production services become widely available.

F2: Carbon and Resource Aware Federation

> **Recommendation:** Take the output of the ongoing NetDRIVE project as an input to federation governance; adopt an agreed sustainability vocabulary and embed carbon- and resource-aware decision making into federated computing; and deploy a coordinated training and knowledge-exchange programme covering carbon-aware computing and sustainability assessment.

Federation creates opportunities to make resource use and environmental impact more visible across services and providers. Shared approaches can help users, providers, and funders understand the consequences of hardware, software, and operational choices more consistently. As approaches to carbon-aware computing mature, federation could provide a common basis for incorporating sustainability considerations into user practice, software and workflow design, resource management, and infrastructure investment decisions.



Chapter 5

One federation, two roadmaps: Shared foundations and coordination

Chapters 1 to 4 have established the federation vision, evidence base, priorities and recommendations. In the brief for the NFCS Programme, UKRI requested separate roadmaps for compute federation and data services federation. These represent complementary delivery pathways within a single federation programme.

As described in §2.3, organisations may participate in federations at different levels according to their needs and circumstances. For that participation to result in coherent federation, a number of shared capabilities must be developed and maintained collectively. These shared foundations span compute and data services and represent the aspects on which both roadmaps depend.

Federation is expected to grow progressively through the voluntary adoption of common approaches by providers and research communities, alongside federation-by-design requirements in future investments. As participation expands, maintaining consistency in standards definitions and interoperability arrangements becomes increasingly important.

This chapter describes the shared foundations, dependencies, coordination mechanisms, and assumptions that connect them before the roadmap options are presented in Chapters 6 and 7. Figure 5.1 previews how these pieces fit together, before each is introduced in detail below.

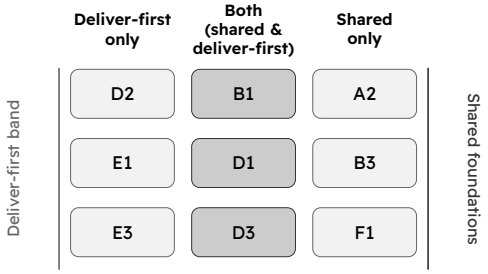


Figure 5.1: One federation, two roadmaps. The deliver-first band identifies what should start first; the shared foundations identify what should be built once for both compute and data services. The two roadmaps therefore follow different delivery pathways while remaining coupled through common foundations and coordination mechanisms.

5.1 Shared foundations and their dependencies

Although compute and data services are described in two roadmaps, they are complementary delivery pathways within a single federation programme. The shared foundations represent those aspects that both roadmaps depend upon. Together they

provide the common capabilities that enable a coherent federation and establish the conditions for interoperability across the UKRI DRI landscape. This is summarised in Table 5.1.

Umbrella	Why both roadmaps depend on it
B1	Shared approaches to data movement, discovery, and access across federated services.
D1, A2	Trusted authentication, authorisation, and project-based access across providers.
D3, F1	Shared approaches to stewardship, allocation, trust, and resource management.
B3	Trusted operation of sensitive, commercial, and regulated data environments.

Table 5.1: The four foundations both roadmaps depend on.

B1: Data interoperability Federated compute workloads depend on the ability to move, access, and discover data across organisational boundaries. Data federation depends on the same standards for interoperability, metadata, and access. Common approaches to data movement and access are therefore foundational to both roadmaps.

D1 and A2: Identity and access Both roadmaps require a consistent approach to identity, authentication, and authorisation. Researchers should be able to access federated resources using recognised identities and appropriate project authority, regardless of whether they are accessing compute services, data services, or both.

D3 and F1: Governance, accountability and allocation Federation requires harmonised approaches to governance, accountability, and allocation. Questions of responsibility, trust, and resource management span both compute and data services and cannot be managed independently within either roadmap.

B3 Secure environments Working with sensitive, commercial, and regulated data requires secure operating environments that can function across a federated infrastructure. Shared approaches to secure storage, trusted research environments, and controlled access with compatible security patterns are therefore shared requirements for both compute and data federation.

Together, these represent the shared investment layer of federation. They are delivered once but serve both roadmaps, and that is the test for inclusion here: neither roadmap can deliver them alone, so each is built once and serves both.

These four are therefore not the same set as the deliver-first band identified in §4.1, and the two answer different questions. The deliver-first band answers what must *start* first, on priority and urgency. The shared foundations answer what must be



built once because neither roadmap can deliver it alone. The memberships differ in both directions: **D2**, **E1** and **E3** start first but are each delivered within one roadmap and relied on by the other; **A2**, **B3** and **F1** are built once for both without being urgent enough to lead.

Dependency asymmetry Both roadmaps depend upon the same shared foundations, but they depend on them in different ways and at different stages of delivery. Data federation depends on identity, governance, and contractual frameworks from the outset — ownership, custody, access permissions and retention must be resolved before services can operate as part of a wider federation. Compute federation depends on the same foundations, but can progress further through standards development and testing before large-scale adoption, since capabilities such as workflow portability can mature while governance and identity arrangements continue to develop. The implication is not that governance and identity are less important for compute federation, but that standards development, testing environments, and enabling activities can proceed in parallel, reducing risk while ensuring later capabilities are built on stable foundations.

5.2 Maintaining coherence across the roadmaps

The shared foundations will not emerge automatically through independent activity. Mechanisms are needed both to develop the foundations in the first place and to maintain them as federation grows. Three complementary mechanisms are proposed: federation working groups, testing environments, and a federation coordination function. These roles are summarised in Table 5.2.

Mechanism	Purpose
Working groups	Maintain standards and stewardship of federation foundations.
Testing environments	Demonstrate interoperability, reduce deployment risk, and facilitate on-boarding.
Coordination Function	Convene stakeholders, broker standards, and oversee delivery of federation.

Table 5.2: Mechanisms to maintain coherence across the developing federation.

Federation working groups The shared foundations described in Section 5.1 require on-going development and stewardship if they are to provide coherence across providers, research communities and funders. Standards, governance, and conformance criteria will need to evolve as new services participate, technologies mature, and user requirements change. The working groups provide the mechanisms through which standards, reference architectures, governance arrangements, and operating practices can be developed collaboratively and maintained over time.



Four cross-cutting working groups are proposed aligned to the four shared foundations.

Data interoperability is responsible for **B1**. Identity and access is responsible for **D1** together with **A2**. Governance and accounting is responsible for **D3** with **F1**. Secure environments is responsible for **B3** and the associated sensitive-data patterns.

Testing environments Many of the questions raised by the programme cannot be resolved through consultation and standards development alone. Testing environments provide a mechanism through which federation capabilities can be developed, evaluated, and verified in practice before they are adopted at scale.

The testing environments serve three complementary purposes: First, they close knowledge gaps by letting questions be settled empirically, providing evidence about what works in practice. Second, they provide a means of verifying conformance to the shared foundations. Third, they reduce delivery risk by allowing new capabilities to be tested before deployment at scale.

A recurring theme across the programme was the need for shared environments in which providers and research communities could experiment with federation approaches before committing to operational deployment. Examples include: data-transfer test-beds, identity integration test-beds, and secure-environment test platforms.

The barrier today, particularly for newcomers and less-established research communities, is not appetite but the lack of understanding and resources to try something out at all. Environments spanning the technology, community, and governance pillars let people explore federation capabilities without building them from scratch.

The testing environments are an enabling capability for federation, not a federation service in their own right. Their purpose is to develop evidence, verify interoperability, and reduce adoption risks.

Federation Coordination The shared foundations, working groups, and testing environments provide a mechanism through which federation can be developed. However, they don't directly provide coordination across the wider federation programme. There needs to be a common point through which priorities can be aligned, dependencies managed, and adoption supported.

A federation coordination function is needed to convene the stakeholders, support the working groups, oversee the testing environments, promote the adoption of federation capabilities, and ensure consistency across the developing federation. This becomes increasingly important as new research communities engage at different levels of federation maturity, and as new investments are commissioned against federation requirements from the outset. The coordination function does not operate



services on behalf of participating providers; its role is to support collaboration, stewardship and alignment across a diverse and independently operated landscape.

NFCS has demonstrated many of the capabilities required of such a function through the delivery of 26 flexible-funded projects, the formation of advisory and community boards, and a sustained programme of community engagement. The precise form of any future coordination function remains a matter for funders and stakeholders to determine.

5.3 Delivery principles and assumptions

The preceding sections identify the shared foundations of federations and the mechanisms required to develop and maintain them. The roadmap options presented in Chapters 6 and 7 explore the different delivery pathways, but all depend on a common set of principles and planning assumptions. These are described below and provide a consistent basis for comparing, appraising, and challenging the options presented in subsequent chapters.

Delivery in both roadmaps is incremental and evidence-led: capabilities should be prototyped, proven in the testing environments, and scaled once they work, so that investment can be committed in tranches against demonstrated progress rather than as a single unbounded commitment.

That is what the decision gates in each roadmap are for. Increments are not, however, a licence to fund foundations partially: each roadmap names the element that must be right before services come to depend on it and cannot be retrofitted afterwards.

What these chapters do not contain Neither roadmap is a delivery plan. Risk registers, key performance indicators, governance structure, and operational procedures depend on which options are ultimately selected and would be developed as part of subsequent delivery planning. Risk is appraised here at the level of the option analysis; the measures that would populate a KPI framework are the adoption thresholds attached to the decision gates.

Planning assumptions These assumptions are stated explicitly so that timings, costs and delivery approaches can be reviewed and challenged independently as the roadmap develops.

Table 5.3: Assumptions underlying the roadmap chapters, cited by ID from the option, appraisal and phase tables.

ID	Assumption	Basis
Timing		

continued overleaf



Assumptions, continued

ID	Assumption	Basis
asm-1	Funding commences Q2 2027; dates rebase from the actual start, though sequence and dependencies hold regardless. Later phases are indicative, not commitments; enablers run throughout.	No committed funding at writing; quarter, not month, precision is honest here. Per July 2026 consultation: *deliver later does not mean start later.*
Cost basis		
asm-2	Effort is in FTE-years, build and run shown separately; cash figures aren't carried forward despite appearing in some reports. Convening/secretariat are recurrent, at roughly four to six events a year.	Reports' cash bases vary (price year, grade, overheads, capital), so a consistent effort figure is more usable. Recurrent obligations do not taper; conflating with build cost loses the data services argument.
asm-3	Only the preferred option is costed; alternatives carry a relative indication of what a lower figure postpones. Phase effort follows strawman bands (S = 1–2, M = 3–6, L = 8–15 FTE-years plus coordination), cross-checked against, not summed from, project reports.	Conversion is a proposal for challenge. Reports' estimates are scoped to recommendations, not sequenced options, so summing double-counts **asm-5**'s shared mechanisms; two reports were also unavailable.
asm-4	Twelve sites are in scope, six early adopters in phase 1, from at least three council domains; site count drives **B1**/**D1** scaling and the gate thresholds.	Placeholder pending provider confirmation. Domain spread matters: sites concentrated in mature communities would pass without testing cross-maturity federation.
asm-5	Working groups and testing environments are single shared instances for both roadmaps, costed once in compute (roughly half attributable to data); not to be summed across chapters.	Same instances, not parallel ones; roadmap-specific work is costed separately where required.
asm-6	Storage capacity cost is not estimated.	Unit costs and demand figures are absent from the programme's evidence base — recorded as **sto-g2**
Scope and method		
asm-7	Umbrella priority and urgency are carried unchanged from strawman, which is frozen for this document.	Final project outputs are reflected in the evidence chapters, not by re-cutting the strawman
asm-8	The preferred option is offered for testing, not a settled conclusion; every figure in the option, appraisal and phase tables is a first-pass estimate owned by its chapter.	Consultation reshaped placements between strawman proposals; the same applies here, so collaborators iterate rather than treat blanks as unassigned.
asm-9	Building on the existing base (§5), conformance is voluntary, applying at the point of interoperation for existing services; future investments may incorporate federation requirements via funding, procurement or service specifications.	Neither NFCS nor any single research council has authority to mandate change; the instrument available is specification



Chapter 6

Compute federation: Roadmap and pathways

This chapter sets out alternative pathways for delivering compute federation across the UK digital research infrastructure landscape. The preceding chapters established the case for federation, identified the priority recommendations arising from the NFCS evidence base, and described the shared foundations required across both the compute and data-services roadmaps. This chapter focusses specifically on how compute federation could be implemented.

This is fundamentally a challenge of interoperability between independently operated services. The roadmap therefore focusses on the standards, governance, identity infrastructure and operational coordination needed to enable services to participate in a wider federation and offer more discoverable and accessible services for users.

There is more than one way to reach this outcome. The evidence suggests several plausible implementation approaches that differ in where investment is initially focused, what capabilities are prioritised, and how rapidly federation is expanded. These options are examined below before a preferred pathway is identified.

The preferred pathway does not represent a final decision, but is presented as a proposition arising from the NFCS evidence base for discussion with the community, the boards, and UKRI. All options assume the shared foundations and coordination mechanisms described in Chapter 5. The primary question considered here is not whether federation should occur, but how it should be delivered. For compute federation, the element that must be settled before services come to depend on it, and cannot be applied retrospectively, is federated identity together with the project-authority model behind it (**D1**): services that onboard users without trusted project context accumulate local access arrangements in its place, and those cannot be unpicked once users and workloads depend on them.

6.1 Implementation options

The delivery options outlined below differ in their centre of gravity. That is, what is funded first and what follows, and what is deferred. These are not mutually exclusive packages of work, each contains elements of the others - they differ in focus, pace, and sequencing. A do-nothing baseline is also considered. They should be understood as alternative responses to the priority $\times$ urgency analysis presented in Chapter 3. That analysis identified a set of deliver-first capabilities on which many other federation capabilities depend. The principal difference between the options is therefore how they treat those dependencies.

Only the preferred option is costed (**asm-2**, **asm-3**); the others carry a relative indication, with a note of what a lower figure postpones rather than avoids. Effort is in FTE-years, with non-FTE costs shown separately. Which of these options pair coherently with the data services options is set out in §5.2.

Option 1: Foundation-led Federation (Preferred)

This option prioritises the common foundations before introducing more advanced or user-facing capabilities (Table 6.1). It is based on the view that sustainable federation depends first on the trusted approaches to identity, governance, cybersecurity and interoperability. The objective is to establish the mechanisms through which independently operated services can participate in a coherent federation and create a stable platform on which later capabilities can be adopted with confidence. This aligns well with the dependency structure identified in the evidence base, reducing the risk of fragmented or incompatible delivery. The trade-off is that some user-visible benefits arrive later than under more front-loaded approaches; however, the resulting federation is more sustainable and better able to support long-term growth across diverse providers and research communities.

Option 2: Standards-led federation

This option prioritises the definition of technical standards and conformance frameworks as the primary mechanism for federation. It is based on the view that interoperability can emerge bottom-up through the widespread adoption of common approaches and standards, allowing individual participants to determine their own pace and approach to implementation.

The principal advantage is that it minimises initial investment and reduces central coordination overhead; but risks producing technical compatibility without shared governance, identity structures, and operational processes required for coherent federation.

Option 3: User-experience led federation

This option prioritises visible improvements to the researcher experience from the outset, focussing at first on the single front door, on-boarding, and service discovery (Table 6.1). The main advantage is that researchers see tangible benefits earlier than more foundation-oriented approaches. The trade-off is that user-facing benefits may be introduced before the underlying federation required to support it is fully established. This risks improving how users encounter federation without fully addressing the organisation and technical fragmentation that federation is intended to overcome.

Option 4: Comprehensive federation

This option prioritises shared foundations, standards, user-facing services, and advanced capabilities concurrently as one coordinated programme (Table 6.1). It is based on the view that the fastest route to a mature federation is to develop all major capabilities in parallel, not sequentially.



The principal advantage is speed, allowing a broad range of federation capability and user benefits to be delivered in a shorter timeframe. The trade-off is increased complexity, cost, and delivery risk. It reduces the flexibility to learn from earlier phases of delivery and increases the consequences of decisions made before the shared foundations have fully matured.

Option 5: No coordinated federation

This counterfactual assumes no dedicated federation programme is established, and existing arrangements continue to evolve independently (Table 6.1). It reflects a view that federation should emerge organically through existing investments, bilateral collaborations, and voluntary adoption of common practices, without coordinated national activity. The key advantage is that it avoids the need for dedicated federation funding and imposes no additional coordination requirements on providers. The trade-off is that the challenges identified in the evidence base remain largely unresolved. Progress depends on individual initiatives aligning by coincidence rather than by design with increased risk of duplicated effort and fragmentation, and a missed opportunity to realise federation's benefits at national scale.

6.2 Appraisal

Options are appraised against seven criteria, weighted towards delivery, reflecting the principal concerns raised throughout the evidence base. The scoring is qualitative; the evidence base does not support numeric scoring and the implied precision would be misleading.

In the two *risk* and *cost* rows, a low rating is favourable. Disagreement is expected to concern the weighting of criteria, not the scores themselves.

Production adoption means in production at provider sites, not demonstrated in a testing environment.

The appraisal exposes the trade-offs between the alternatives; it does not identify a single objectively correct option. It is summarised in Table 6.2.

6.3 The preferred pathway

Why this option *Foundation-led* is the only option that aligns fully with the dependency structure identified in the evidence base and the deliver-first priorities established in Chapters 3 and 4. It scores well on dependency integrity and stranded-investment risk at a moderate cost in time to first benefit. This conclusion is also consistent with the programme's strongest areas of evidence, **D3** with 43 recommendations, **B1** with 35 and **D2** with 32.

Critically, this option supplies the governance, identity, and trust arrangements that allow federation standards to be consistently adopted. By establishing these foundations first conformance becomes a credible route for existing services to join in with federation, while future investments could incorporate federation requirements from the outset.



Option	Leads with / defers	Reliance on shared mechanisms	Indicative cost	Confidence
Foundation-led *(preferred)*	*Leads:* **D1** and **D3** as shared national infrastructure; **D2** as compute-specific governance infrastructure; **B1** and **A2** as conformance standards. *Follows in phase 1:* **E1**, **E3**, once an authority model exists to sit behind them, scoped across services rather than to a single one. *Defers:* **A1**, **A3**, **C1**, **C2a**, **C2b**, **F1**. *Horizon:* **C3**, **F2**. Enablers **A2**, **B3**, **E2** throughout.	High — WGs own the standards; conformance testing gates adoption	80 FTE-years over three phases, plus recurrent convening and testing-environment operation	High
Standards-led	*Leads:* **B1**, **A2** specifications and a conformance regime. *Defers:* everything else to provider pace; **D1–D3** become voluntary guidance.	Highest — WGs and conformance testing are the only delivery vehicle	Lower up front, perhaps a third; defers cost without avoiding it, as coordination persists and adoption stretches	Moderate
User-experience led	*Leads:* **E1**, **E3**. *Back-fills:* **D1–D3**, **B1** behind a live user-facing route.	Low — and this is part of the case against it	Comparable overall; front-loaded, with rework exposure once the authority model arrives	Moderate
Comprehensive	All 17 umbrellas resourced concurrently as one programme. Included for completeness and coverage should the funding envelope allow, not as a recommended starting point.	Moderate — coupling still needs owners, but sequencing pressure is lower	Higher, perhaps half again; highest coverage per unit time	Low
Baseline	Existing federations and bilateral arrangements continue, unchanged and uncoordinated.	None	Nil committed; cost accrues as duplicated effort at institutions and providers	High

Table 6.1: The map: appraised options for the pathway to compute federation. Confidence refers to the evidence behind the option, not to the cost figure. First-pass estimates for challenge (**asm-8**).

Why user-experience led is not recommended User-experience led federation prioritises visible benefits. While attractive in principle, the evidence suggests that many of the barriers researchers experience arise not from the absence of user interfaces, but from the weak connection between identity, project authority, access management, and service delivery. Improving the user journey without addressing the underlying foundation risks making federation easier to navigate without making it materially easier to use.

A single front door can only simplify access if the authority model behind it is already



Criterion	Foundation-led	Standards-led	User-experience led	Comprehensive	Baseline
Production adoption	High	Low	Mod.	High	None
Dependency integrity	High	Mod.	Low	Mod.	—
Time to first user-visible benefit	Mod.	Slow	Fast	Fast	None
Risk of stranded investment	Low	Mod.	High	Mod.	—
Evidence strength	High	High	Mod.	Mod.	—
Recurrent-cost exposure	Mod.	Low	Mod.	High	Low
Alignment with 2×2	High	Mod.	Low	Mod.	None

Table 6.2: Appraisal of options against stated criteria. A low rating is favourable in the risk and recurrent-cost rows. First-pass assessment for challenge (**asm-8**).

trusted and understood by participating providers. Delivering user-facing capabilities ahead of governance, identity, and operational frameworks creates significant risk of later rework. The front door remains a phase 1 deliverable under the preferred option, sequenced after the authority model.

Why standards-led is not recommended Standards-led federation prioritises interoperability specifications and conformance frameworks as the primary mechanism for federation. Common standards and approaches are essential for services to interoperate and the evidence base strongly supports their development. However, they do not provide the governance and coordination framework needed to make federation work in practice.

The shared foundations identified in Chapter 5 are intended to be developed and maintained collectively. Standards-led federation risks creating technical compatibility without operational federation, slowing adoption and reducing confidence that a coherent, interoperable federation can be achieved.

The preferred pathway treats standards as a critical component of federation, but places them within a wider framework of governance and coordination frameworks where they can support practical routes for participation in federation.

Why comprehensive federation is not recommended Comprehensive federation offers the faster route to capability coverage, but commits substantial effort and funding before the roadmap's key uncertainties have been reduced. This weakens the evidence-led delivery by reducing opportunities to learn through testing environments and decision gates before later commitments are made. The preferred pathway delivers many of the same long-term outcomes while reducing risk.

What lands when In the first twelve months the testing environments are in operation and producing referenceable artefacts — a conformance suite and an identity integration testbed — while the foundations are still being built, and the first increments of **D1**, **D2** and **E1** are live at early-adopter sites. By the end of phase 1 a researcher signs in once and is recognised with project authority across participating



	Phase 1 Q2 2027 – Q1 2029	**Phase 2** Q2 2029 – Q1 2031	**Phase 3** Q2 2031 onwards
Umbrellas in scope	**D1**, **D2**, **D3**; **B1** and **A2** specifications; **E1**, **E3**	**A1**, **C1**, **C2a**, **C2b**, **F1**, **A3**	**C3**, **F2**; review of deferred items
What a user can do by the end	Sign in once with an institutional identity and be recognised, with project authority, at participating sites across more than one council domain; find what they can access and get support through one route	Submit a workflow that runs across sites without rewriting it; use governed AI research services; see allocation and usage accounted consistently	Route suitable work to quantum-integrated resources; see the carbon and resource cost of where work runs
Depends on	—	**D1**–**D3** in production; **B1** adopted at the phase 1 threshold	Phase 2 gate; maturity of the horizon technologies
Cost: FTE-years	38	30	12
Cost: non-FTE	Testing-environment build and run; 4–6 convening events a year; proof-of-concept hardening; secrets-management licensing; training	Testing-environment run; convening; training	Testing-environment run; convening
Decision gate	**D1** and **D2** in production at 6 sites and the **B1** conformance suite passing at 3, before phase 2 is released	**A1** in production at 6 sites and **F1** accounting live, before phase 3 is released	Horizon review: whether **C3** and **F2** are ready to move into delivery

Table 6.3: The preferred pathway in phases. Dates assume funding commencing Q2 2027 and rebase from the actual start; the sequence and dependencies hold regardless. Later phases are indicative shape, not commitments, and are re-scoped at each gate; enablers run throughout rather than within a phase (**asm-1**). Adoption thresholds follow **asm-4**; effort follows **asm-3**.

sites. Phase 2 makes work portable between them. Phase 3 revisits what was deferred. Nothing here waits five years for its first user.

6.4 Delivery and verification

A key objective of the preferred pathway is to close the knowledge gaps identified in Chapter 3 while delivering federation capabilities into production. The working groups and testing environments described in §5.2 provide the mechanisms through which this is achieved. They are shared instances serving both roadmaps and are costed once (**asm-5**).

The preferred pathway is structured so that each major uncertainty is addressed through either experimentation, implementation, or governance activity before later phases depend on the result.

The identity integration testbed addresses **fed-g2**, since which authority models institutions can actually implement is not answerable from the written evidence. The conformance suite addresses **fed-g3**, establishing what passing the data-movement standard means in practice before it is required of providers. Phase 2 orchestration



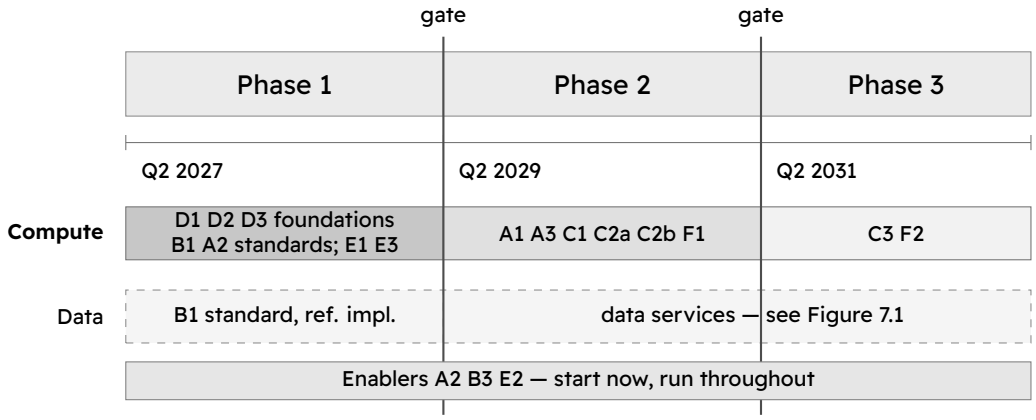


Figure 6.1: The preferred pathway drawn across the map: phases, the continuous enabler band, and the decision gates at which later phases are re-scoped. The data services track is shown for coupling; it is foregrounded in the data services chapter, where this figure appears with the emphasis reversed.

work addresses **fed-g4** on portability across heterogeneous schedulers. **Fed-g1** and **fed-g5** are closed by the working groups, being questions of definition and operating model, not implementation. **Fed-g6** is addressed by the domain spread required of the phase 1 early adopters (**asm-4**), which tests federating across maturity levels rather than within the most mature research communities.

The decision gates ensure that key uncertainties are resolved and that evidence is generated before additional capability is deployed or further investment committed.

6.5 Implications for future DRI funding activities

The five propositions below share a single argument: federation's foundations carry recurrent obligations that do not end when a grant does, and capabilities funded as time-limited projects lapse precisely when adoption, and the obligations it brings, begin in earnest. This is the same logic behind the data roadmap's proposition to fund retention as a recurrent obligation instead of a project output (§7.5), and behind the cost-recovery model recorded as a knowledge gap in both roadmaps (**fed-g5**, **sto-g3**). Financial sustainability is not separate from the propositions below: it is what each is an instance of.

Fund the foundations as infrastructure, not as project outputs. Identity, cybersecurity, and the operating model carry recurrent obligations that do not end with a grant. These capabilities must be maintained throughout the life of the federation and therefore require funding models aligned with infrastructure operation rather than short-lived project delivery.

Require federation-by-design in future DRI investments as a condition of commissioning new services. Federation requirements are easiest to introduce at commis-

sioning or procurement. Once services are operational integration becomes progressively harder and more expensive to retrofit.

Fund the shared testing environments as a standing capability. They close knowledge gaps, verify conformance, and de-risk increments; funded as a project, they lapse precisely when conformance begins to matter.

Separate recurrent operating cost from capability build. Federation combines one-off capability build with long-term operation obligations. Treating them both as the same kind of expenditure obscures the true cost of federation and creates sustainability risks once adoption begins.

Fund a federation coordination function as shared national infrastructure. The roadmap assumes a coordination capability that can convene stakeholders, support the working groups, oversee conformance, support adoption across providers and research communities, and maintain coherence across the federation as it grows.





Chapter 7

Data services for compute federation: Roadmap and pathways

This chapter sets out alternative pathways for delivering the data services on which compute federation depends. The preceding chapters established the case for federation, identified the priority recommendations arising from the NFCS evidence base, and described the shared foundations required across both the compute and data-services roadmaps. This chapter focusses specifically on how federated data services could be implemented.

Its scope is the data services the federation depends on. The wider storage needs of UK computational research are addressed where the programme's evidence reaches them and recorded as a knowledge gap where it does not (**sto-g1**). As with compute, this is fundamentally a challenge of interoperability between independently governed services. Data custody is irreducibly distributed across institutions, national facilities, commercial cloud, and discipline repositories, so the roadmap focusses on the standards, metadata, and security patterns through which those holdings can be moved, discovered and used across the federation.

There is more than one way to reach this outcome. The evidence suggests several plausible implementation approaches that differ in where investment is initially focused, what capabilities are prioritised and how rapidly provision is expanded. These options are examined below before a preferred pathway is identified.

The preferred pathway does not represent a final decision, but is presented as a proposition arising from the NFCS evidence base for discussion with the community, the boards, and UKRI (**asm-8**). All options assume the shared foundations and coordination mechanisms described in Chapter 5. For data services the element that must be settled before services come to depend on it, and cannot be applied retrospectively, is data contracts and custody (**D3**): provenance, permissions and retention obligations are fixed when data is created, not when a governance framework is agreed. Nothing counts as delivered while it remains in a testing environment: the measure is adoption in production by DRI service providers and research communities.

7.1 Implementation options

As with the compute roadmap the delivery options outlined below differ in: what is funded first and what follows, and what is deferred. These are not mutually exclusive packages of work. Each contains elements of the others - they differ in where they place the initial emphasis. Unlike compute federation, where the primary challenge is

establishing a shared authority model across independently operated services, data federation must contend with irreducibly distributed custody arrangements spanning institutions. Both roadmaps assume the shared foundations and coordination mechanisms described in Chapter 5, but use them differently. The principal question is therefore whether federation should focus first on interoperability, capability provision, evidence gathering, or some combinations of all three. These alternatives are explored below.

A do-nothing baseline is also considered, which for data services is materially worse than for compute: capacity cost accrues whether or not it is planned for, and custody obligations attach to data as it is created.

Only the preferred option is costed (**asm-2**, **asm-3**); the others carry a relative indication, with a note of what a lower figure postpones without avoiding. Effort is in FTE-years, with non-FTE costs shown separately. Storage capacity cost is not estimated: the unit costs and demand figures required to size it are absent from the programme's evidence base (**asm-6**, **sto-g2**), which is itself among the findings this roadmap carries forward.

Option 1: Standards-led data services (Preferred)

This option prioritises interoperability across independently governed data services. It is based on the view that data custody, stewardship, and provision will remain distributed across institutions, making interoperability more important than consolidation. The objective is to establish common approaches to data movement, metadata, discovery, access, and security that allow existing and future services to participate in federation. The principal advantage is that it aligns with the strongest areas of evidence and the realities of the existing landscape. It allows services to interoperate without requiring changes in ownership, custody, or operating responsibility. The trade-off is that it does not itself address future storage-capacity requirements. Instead capacity decisions are informed by a dedicated demand measurement track, reducing the risk of committing significant resources before future demand and costs are properly understood.

Option 2: Nationally provided data service

This option prioritises centralised provision of federated storage and capacity. It creates a nationally provided data service as the primary route to federated data capability, addressing interoperability with existing provision as a later step (Table 7.1). It is based on the view that a single provider can simplify governance, user access, support, and service delivery while reducing the need to broker agreements across multiple organisations. The principal advantage is simplicity: a single provider offers a clearer operational model and more consistent service with fewer coordination points than a fully federated landscape of independently governed services. The trade-off is that much of the existing data landscape already exists outside any single provider and carries obligations that cannot be readily transferred. As a result, the option risks



concentrating on provision while deferring the interoperability challenge that federation is intended to address. It also introduces significant uncertainty around scale: without robust evidence on demand and costs, any nationally provided service would need to be sized before the programme has established what capacity was required.

Option 3: Demand-led, measurement first

This option prioritises measuring demand, volumes, growth rates and unit costs before committing to a delivery pathway. It is based on the view that future demand and costs should be understood before major decisions are taken about standards, services, or capacity provision. The objective is to reduce the risk of over- or under-investment by ensuring that future decisions are informed by a clearer understanding of the scale and cost of the problem. The principal advantage is that it directly addresses the programme's most significant evidence gaps, particularly around demand and storage economics. Capacity decisions taken after such measurements are better informed and easier to justify. The trade-off is delay. Deferring interoperability standards and federation mechanisms until the measurement activity is complete postpones benefits to both the data and compute roadmaps.

Option 4: Comprehensive data provision

This option develops standards, national provision, and capacity together as one coordinated programme. It is based on the view that the fastest route to mature data federation is to address all these elements together instead of sequencing them over time. The principal advantage is speed, allowing a broad range of federated data capabilities to be delivered within a shorter overall timeframe. It aims to accelerate adoption and provide earlier access to a comprehensive federated data environment.

The trade-off is increased complexity, cost and delivery risk. Significant decisions about capacity provision are made before the programme's principal uncertainties around demand, cost, and long-term retention requirements have been resolved. As a result, the option carries the greatest risk of stranded investment and reduces the opportunity to learn from evidence generated through phased delivery.

Option 5: No coordinated data federation

This option is provided as a counterfactual: no dedicated programme is established, and existing arrangements continue to evolve independently (Table 7.1). It reflects the view that interoperability and coordination should emerge organically through voluntary contributions, not through coordinated national activity.

As with the equivalent option for the compute roadmap, this avoids the need for dedicated federation funding and imposes no additional coordination obligation on participants in federation. The trade-off is that the challenges identified throughout the evidence base remain largely unresolved. Unlike the equivalent compute baseline, it is not cost-free: capacity and custody costs accrue regardless, contracting effort continues to be duplicated at every institution, and data generated without an agreed custody arrangement cannot be brought into a federation retrospectively. It



also provides no structured mechanism for addressing the programme's identified knowledge gaps or coordination for future investment decisions.

Option	Leads with / defers	Reliance on shared mechanisms	Indicative cost	Confidence
Standards-led *(preferred)*	*Leads:* **B1** — specification, reference architecture, and reference implementation — with metadata, cataloguing, and discovery as a named strand; **B3** patterns for sensitive data. *Funded tracks alongside:* demand measurement, and targeted implementation where the evidence is already mature. *Defers:* **F1** capacity accounting to phase 2. *Horizon:* **F2**. Enablers **A2**, **B3**, **E2** throughout.	High — WGs own the standards; conformance testing gates adoption	48 FTE-years over three phases, plus a commissioned demand study. Capacity excluded (**asm-6**)	High
National data service	*Leads:* nationally provided federated storage and capacity. *Defers:* interoperability with existing institutional, cloud and discipline provision to a later integration step.	Moderate — a single provider reduces the need to broker between many	Higher, and capital-weighted; cannot be sized at all without the demand figures the evidence base lacks	Low
Measure demand first	*Leads:* demand, volume, growth, and unit-cost measurement across UK computational research. *Defers:* the standard and all provision until the measurement reports.	Low in the first phase; rises once a route is chosen	Lower up front, perhaps a fifth; defers rather than avoids cost, and delays the standard compute federation depends on	Moderate
Full provision	Standards, national provision and capacity funded concurrently. Included for completeness and coverage should the funding envelope allow, not as a recommended starting point.	Moderate — coupling still needs owners, but sequencing pressure is lower	Highest; highest coverage per unit time	Low
Baseline	Existing arrangements continue, unchanged and uncoordinated.	None	Nil committed; capacity and custody costs accrue regardless, unmanaged	High

Table 7.1: The map: appraised options for the pathway to federated data services. Confidence refers to the evidence behind the option, not to the cost figure. First-pass estimates for challenge (**asm-8**).

7.2 Appraisal

Options are appraised against seven criteria, weighted towards delivery and informed by the principal concerns raised in the evidence base. For data services, particular weight is placed on dependency integrity and the management of uncertainty around



future demand, capacity requirements, and long-term stewardship obligations.

The scoring is qualitative; the evidence base does not support numeric scoring and the implied precision would be misleading.

In the two risk and cost rows, a low rating is favourable. Disagreement is expected to concern the weighting of criteria, not the scores themselves. That is what makes this table usable in board and workshop discussion.

Production adopted means in production at provider sites, not demonstrated in a testing environment.

The appraisal exposes the trade-offs between the alternatives, Table 7.2, instead of identifying a single objectively correct option.

Criterion	Standards-led	National service	Measure first	Full provision	Baseline
Production adopted	Mod.	Mod.	Low	High	None
Dependency integrity	High	Mod.	Low	Mod.	—
Time to first user-visible benefit	Mod.	Slow	Slow	Fast	None
Risk of stranded investment	Low	High	Low	High	—
Evidence strength	High	Low	High	Low	—
Recurrent-cost exposure	Mod.	High	Low	High	High
Alignment with 2×2	High	Low	Mod.	Mod.	None

Table 7.2: Appraisal of options against stated criteria. A low rating is favourable in the risk and recurrent-cost rows. The baseline carries high recurrent-cost exposure because capacity and custody costs accrue whether or not they are planned for. First-pass assessment for challenge (**asm-8**).

7.3 The preferred pathway

Why this option Standards-led data services is the only option that aligns fully with both the distributed nature of data custody and the strongest areas of evidence identified by the programme. Data holdings already exist across institutions often with governance, retention, and contractual obligations that cannot be transferred elsewhere. Federation therefore depends first on establishing common approaches.

This conclusion is consistent with the programme's strongest areas of evidence, particularly **B1** (35 recommendations) and **B3** (12 recommendations). The programme produces a web API specification, a documented reference architecture, an open-source reference implementation, and a recurring recommendation that participating services interoperate through common standards.

The trade-off is that this option does not directly address future storage capacity requirements. Instead capacity decisions are informed by a dedicated demand-measurement track running alongside delivery (**sto-g2**). This allows interoperability to advance while reducing the risk of committing significant resources before demand and costs are properly understood.



The compute foundations then provide the authority model that a standards regime alone lacks. Conformance is the mechanism through which existing services can participate in federation, not a requirement imposed universally across the landscape.

Why a national data service is not recommended Consolidating provision into one nationally funded service is the intuitively attractive route and the argument against it is the shape of the existing landscape rather than its cost. The evidence suggests that the principal challenge facing data federation is not the absence of a national provider, but the need for interoperation across data holdings that already exist. Consolidated provision doesn't reduce the need for interoperability, it simply defers the problem federation is intended to solve. The preferred pathway invests first in the standards that allow services to participate in federation irrespective of where the data is ultimately hosted.

Why measurement first is not recommended Measurement is necessary but not sufficient, and it is not free of cost. The evidence base cannot yet size future capacity requirements, but it does provide sufficient evidence to justify the development of interoperability standards. Deferring the data-movement and access standard until a demand study reports would delay both data federation and the compute roadmap. The preferred option takes the same measurement work and funds it as a track running alongside delivery, which secures the evidence without stalling the standard that everything else rests on.

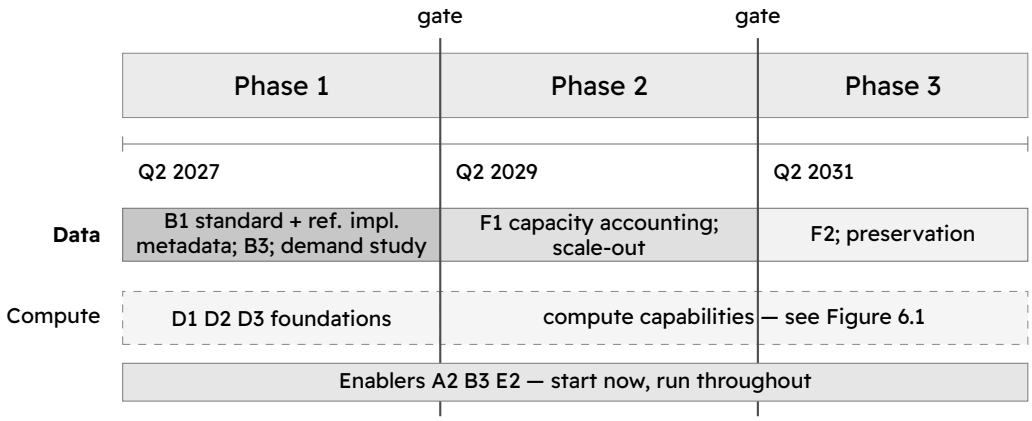


Figure 7.1: The preferred pathway drawn across the map: phases, the continuous enabler band, and the decision gates at which later phases are re-scoped. The compute track is shown for coupling, with the identity and governance foundations in phase 1 on which this roadmap depends; it is foregrounded in the compute federation chapter, where this figure appears with the emphasis reversed.

What lands when In the first twelve months the data-transfer testbed is in operation for users and providers, and the reference implementation is available for sites to test against. The demand study is commissioned in this period, not deferred to a later phase. By the end of phase 1 a researcher moves data between participating sites



	Phase 1 Q2 2027 – Q1 2029	**Phase 2** Q2 2029 – Q1 2031	**Phase 3** Q2 2031 onwards
Umbrellas in scope	**B1** specification, reference implementation and metadata strand; **B3** patterns; demand study	**F1** capacity accounting; scale-out adoption; targeted implementation	**F2**; long-term retention and preservation model
What a user can do by the end	Move a dataset between participating sites with one command and one credential; find datasets through a catalogue; work with sensitive data under documented, tested patterns	See storage allocation and usage accounted alongside compute, and request capacity through the same front door	Know where data will live after the grant ends and what that costs; see the resource cost of storage choices
Depends on	**D1**–**D3** from the compute roadmap	Phase 1 gate; **F1** shared with the compute roadmap	Phase 2 gate; the preservation funding decision
Cost: FTE-years	22	18	8
Cost: non-FTE	Commissioned demand study; testing-environment build and run; proof-of-concept hardening; convening; licensing; training. Capacity excluded (**asm-6**)	Testing-environment run; convening; training	Testing-environment run; convening
Decision gate	Reference implementation in production at 6 sites, **B3** patterns at 3 sensitive-data sites, and the demand study reported, before phase 2 is released	**F1** accounting live at 6 sites and custody and retention contracts in force, before phase 3 is released	Horizon review, and a decision on who funds long-term preservation (**sto-g3**)

Table 7.3: The preferred pathway in phases. Dates assume funding commencing Q2 2027 and rebase from the actual start; the sequence and dependencies hold regardless. Later phases are indicative shape, not commitments, and are re-scoped at each gate; enablers run throughout rather than within a phase (**asm-1**). Adoption thresholds follow **asm-4**; effort follows **asm-3**.

using a common interface and credential, finds it through a catalogue, and can work with sensitive data using documented and tested security patterns. Phase 2 makes storage capacity visible and accountable across the federation. Phase 3 settles what happens to data after the grant ends. The first user-visible change arrives well inside the first phase.

7.4 Delivery and verification

A key objective of the preferred pathway is to close the knowledge gaps identified in Chapter 3 while delivering federation capabilities into production.

The working groups and testing environments described in §5.2 provide the mechanisms through which this is achieved. They are shared instances serving both roadmaps and are costed once (**asm-5**).

The coordinated cybersecurity baseline that the **B3** patterns map to is likewise delivered once, in the compute roadmap (**D2**), and applies across both; it is not phased or



costed again here. What is specific to data services is which knowledge gaps each capability closes, and when. The data-transfer testbed addresses **sto-g5**, since egress and tiering economics across a mixed institutional and cloud estate cannot be settled on paper. The hardening route from the evaluated encrypted-storage proof of concept to a production reference implementation is what makes **B3** adoptable rather than demonstrable. The metadata and cataloguing conformance strand addresses **sto-g4**, where discipline practice diverges most. The commissioned demand study addresses **sto-g2**, and its report is a phase 1 gate condition rather than a deliverable that can slip. **sto-g1** and **sto-g3** fall to the working groups and to funders respectively, being questions of scope and of who pays.

The decision gates therefore ensure that evidence on demand, interoperability, and operational practice is gathered before larger commitments to adoption, accounting, and long-term stewardship are made.

7.5 Implications for future DRI funding activities

This proposition, and the operating and cost-recovery question it raises (**sto-g3**), reflects the same financial-sustainability argument set out in the compute roadmap (§6.5). Five propositions follow, each intended to be liftable into a call specification.

Commission the demand, volume and unit-cost measurement the evidence base lacks. No capacity decision whether national, institutional, or commercial can be sized without evidence on demand, growth rates, and costs. Every option considered in this roadmap is weakened by the absence of these data(**sto-g2**). Measuring these should therefore proceed as a funded activity in parallel with standards development and adoption.

Make conformance to the data-movement and access standard a condition of commissioning new storage provision. As with the federation-by-design in the compute roadmap, interoperability requirements are easiest to introduce at commissioning and difficult to retrofit later. Requiring conformance at procurement ensures that new investments contribute to federation from the outset rather than creating additional integration work in the future.

Fund metadata, cataloguing and discovery explicitly. Federated data that cannot be found is of limited value. Metadata, cataloguing, and discovery are therefore as important as data movement itself, yet the B1/B2 consolidation removes a distinct home for this work. Explicit funding is needed to address the knowledge gap around cross-disciplinary metadata and cataloguing practice (**sto-g4**). Further discussion is found in the Appendix.

Fund retention and preservation as a recurrent obligation, not a project output. Data retention creates obligations that can persist long after the activity that gen-



erated the data has finished. Treating retention and preservation as project outputs risks creating unfunded responsibility once grants end. The unresolved question of who funds those responsibilities remains one of the most significant issues identified by the roadmap (**sto-g3**).

Scope the wider storage needs of UK computational research. Federation and storage provision are related but distinct problems. This programme addressed how independently governed data services can interoperate; it did not establish the overall storage requirements of UK computational research. The wider capacity question there remains open and requires a separate assessment (**sto-g1**).

These propositions complement rather than duplicate national dataset initiatives: for example the National Data Library (24) addresses curated national datasets and their stewardship, while this roadmap addresses how researchers can discover, access, move, and retain research data across a federated infrastructure.



Chapter 8

Roadmap to delivery and next steps

This document is a step in a staged process rather than a finished plan. Chapters 6 and 7 set out evidence-based pathways for compute and data federation, together with the assumptions, dependencies, and knowledge gaps that shape them. This chapter considers what follows: the proposition offered for challenge, the activities required to move from roadmap to delivery, and the arrangements through which federation could be activated and sustained. Decisions on whether and how to proceed remain matters for funders, informed by the providers and research communities who would deliver and use federation.

8.1 Propositions offered for challenge

The roadmaps state a preference plainly so that it can be challenged (**asm-8**). Five things in particular are offered for testing rather than presented as settled:

- The two preferred pathways — *foundation-led* for compute federation, and *standards-led* for data services — and the argument that they are complementary rather than independent choices.
- The appraisal criteria and their weighting. Disagreement is expected to concern how the criteria are weighted rather than the scores themselves, which is what makes the appraisal tables usable in discussion.
- The proposed sequencing and coupling between the compute and data roadmaps, including the role of shared foundations, testing environments, and decision gates.
- The proposal for four cross-cutting working groups aligned to the shared foundations, which is a departure from the July 2026 consultation and is put forward as such.
- The planning assumptions, and every figure that rests on them. None is a costed bid; each is a first-pass estimate owned by the chapter that states it.

This approach has precedent within the programme. The six-theme structure and the priority $\times$ urgency model were put to the flexible-funded projects at the July 2026 consultation, and feedback from that meeting directly reshaped the placement of several umbrella recommendations.

The same spirit applies to the actor-mapping and operational-difficulty tables 8.1 and 8.2 that follow. These reflect first-pass judgements, offered here to be tested and refined rather than presented as settled. Areas of contention or compromise arising from consultation, including one still-open question about how the federation aligns with other live national investments, are recorded in the appendix.

Community refers to organised or representative groups (e.g. community boards, communities of practice); *Users* refers to individual researchers and research teams

Table 8.1: Recommendations mapped to the actors who could plausibly be responsible for delivery. Ticks indicate plausible involvement, not sole ownership; most recommendations require coordination between more than one actor.

Ref	Recommendation	UKRI	Community	New entity (e.g. NFCS-Next)	Users	Infra. providers
A1	Workflow & job orchestration		✓		✓	✓
A2	Federated Access Mechanisms					✓
A3	Industry & public sector access	✓		✓	✓	✓
B1	Data movement, cataloguing, access & storage standard		✓	✓		✓
B3	Federated Trusted Research Environments					✓
C1	Software & container supply chain			✓		✓
C2a	AI research-support services		✓	✓	✓	✓
C2b	Responsible AI governance	✓	✓	✓		
C3	Quantum–HPC integration pathway	✓	✓			✓
D1	Federated identity & access management			✓		✓
D2	Coordinated cybersecurity governance			✓		✓
D3	Federation governance & operating model	✓	✓	✓		
E1	Single front-door, onboarding & support			✓	✓	✓
E2	Skills, RSE capacity & communities of practice		✓		✓	
E3	Culture, widening participation & user needs		✓		✓	
F1	Fair, transparent resource allocation and accounting	✓		✓		✓
F2	Carbon- & resource-aware federation	✓				✓

engaging with services directly. *New entity (e.g. NFCS-Next)* denotes a possible successor coordination function, not necessarily identical in form to NFCS as currently constituted.

Table 8.2: Indicative resourcing (**asm-3**) and operational difficulty per umbrella recommendation. Operational difficulty is a first-pass judgement offered for challenge, not an evidenced figure like the resourcing band.

Ref	Recommendation	Resourcing	Difficulty	Basis / notes
A1	Workflow & job orchestration	L	Medium	Open placement question (K8s) unresolved
A2	Federated Access Mechanisms	M	Low	Single-project origin, contained scope
A3	Industry & public sector access	M	High	Requires external commercial/public-sector engagement beyond the federation itself
B1	Data movement, cataloguing, access & storage standard	L	High	Broadest evidence base (6 projects, 35 recs); absorbed the B2 merge

continued overleaf



Operational difficulty, continued

Ref	Recommendation	Resourcing	Difficulty	Basis / notes
B3	Federated Trusted Research Environments	M	Medium	Technically contained; institutional key-ownership coordination required
C1	Software & container supply chain	M	Medium	Cross-site registry and scanning coordination
C2a	AI research-support services	M	Medium	Emerging area; governance dependency on C2b
C2b	Responsible AI governance	S	Medium	Policy/assurance coordination, though lightly resourced
C3	Quantum–HPC integration pathway	S	Low	Long-horizon, exploratory; low near-term coordination burden
D1	Federated identity & access management	L	High	Highest cross-project convergence (11 projects); requires institution-by-institution adoption
D2	Coordinated cybersecurity governance	L	High	Large scope (32 recs) spanning all providers; single-project evidence base, technically extensive
D3	Federation governance & operating model	M	High	Highest recommendation count (43); spans funding, contracting and governance roles
E1	Single front door, onboarding & support	L	Medium	Wording/scope clarity flagged in consultation; 5-project convergence
E2	Skills, RSE capacity & communities of practice	M	Low	Delivered through existing communities of practice
E3	Culture, widening participation & user needs	M	Medium	Ongoing engagement across diverse and less-engaged communities
F1	Fair, transparent resource allocation and accounting	M	Medium	Scope of mapping flagged as unclear in consultation
F2	Carbon- & resource-aware federation	M	Low	Depends on external NetDRIVE (25) output; contained scope

Operational difficulty reflects how many independent parties must agree or change behaviour, and whether delivery depends on currently unresolved knowledge gaps or open scope questions. It is a first-pass judgement offered for challenge and refinement by the working groups and boards, consistent with the treatment of other planning figures (**asm-8**), and is not derived from the evidence base in the way the resourcing bands are.

8.2 Activating the roadmap

Before delivery can begin, a small number of remaining questions need to be settled because they shape the scale, scope and operating model of any subsequent investment. Table 8.2 highlights a recurring theme throughout the roadmap; the capabilities identified as foundational are also frequently those requiring the greatest coordination across institutions, providers, funders and research communities.

Three of the knowledge gaps recorded in §3.3 are therefore particularly important: what “federated” requires of an existing service in operational terms (**fed-g1**), what



readiness means for a domain starting from close to zero (**fed-g6**), and the operating and cost-recovery model for the shared foundations (**fed-g5**) all bear directly on the scale and shape of any subsequent investment.

The term *federated* is already in wide use, including in national strategy, without a definition against which a service can be assessed. This creates an unusual situation in which investment decisions may depend on concepts that cannot yet be verified in practice. Resolving **fed-g1**, **fed-g5** and **fed-g6** is comparatively inexpensive in effort, but has a disproportionate impact on the precision of subsequent delivery, adoption, and funding decisions. For that reason, these questions are better treated as activation prerequisites rather than as activities deferred into later phases.

The case for continuity

Federation benefits every party to it, in different ways. Funders get coordination in place of duplicated investment; providers extend their reach without losing control of their own infrastructure; research communities get simpler, more equitable access to compute and data capability regardless of where they are based. None of this is realised automatically: it depends on a coordinating body sitting between them, convening agreement, brokering demand against provision, and overseeing delivery of the shared mechanisms that make the system behave as one.

The roadmap points to the need for a continuing coordination function. Table 8.1 identifies such a role as plausible contributor to ten of the seventeen umbrella recommendations, including most of the deliver-first band, not as sole owner but as one of the several actors whose participation is required for delivery. The capabilities this would require — convening, brokering, and stewardship of the shared foundations — are those the programme has already exercised (§5.2).

Community board input on this roadmap raised a further consideration: if this evidence base, these relationships and this convening capacity are not carried forward, they are not simply paused — they are lost. Significant parts of that capability must be recreated by whichever organisation assumes responsibility next. Transition is possible but is not cost free.

This does not determine who should deliver activation — that remains a decision for funders. The key point is that continuity should be treated as the default assumption, with transition requiring a positive case based on the benefits it would deliver. The evidence gathered through the NFCS suggests that continuity may represent a lower risk and faster path from roadmap to delivery than recreating equivalent capability elsewhere.

Proposed operating model

The roadmap already identifies the mechanisms through which federation would be developed and maintained: federation working groups, shared testing environments, and an overseeing coordination function (Chapter 5). Figure 8.1 makes that descrip-



tion operational, setting out how three functions: convening, brokering, delivering, map onto those mechanisms. The figure illustrates how this could operate in practice.

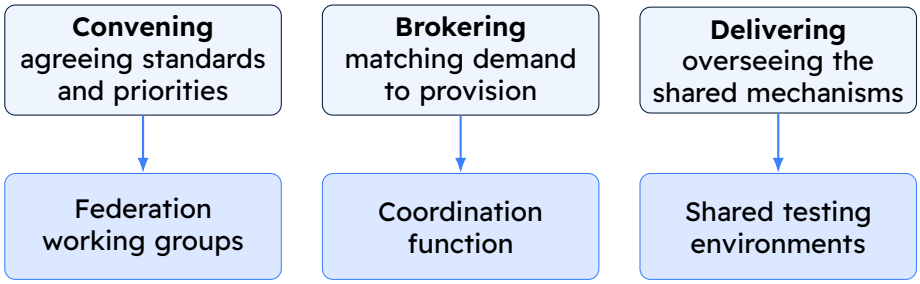


Figure 8.1: The three coordinating functions, mapped onto the roadmap's existing mechanisms.

Convening is delivered initially through the roadmap's four cross-cutting working groups, named for what they do rather than only what they own. Their purpose is to bring funders, providers and research communities together, across research councils and disciplines, to agree standards, priorities, and the shared foundations on which federation depends.

Delivering is the oversight of the shared mechanisms through which federation is developed and adopted, including the testing environments and the shared foundations once they enter production use. The role of the coordination function is not to build and operate every federation component itself but to ensure that agreed standards and implementation activities translate into working capabilities.

Brokering is the coordination function's practical work between funders, providers, and research communities. It translates allocation and investment decisions into operational arrangements; matches community demand against distributed provision, and helps identify gaps where capability, capacity or coordination is missing. A particular benefit of federation is the ability to replace repeated bilateral agreements with common agreements that can be reused across many participants in federation.

The roadmap is deliberately designed to accommodate change. Federation participation is progressive and reversible: services may move to deeper integration, or step back, as their circumstances change (§2.3). Delivery is phased and evidence led, with decision gates that allow later work to be re-scoped in light of observed practice. The horizon band provides a route for incorporating emerging priorities without requiring the roadmap itself to be redesigned. The challenge is not creating a new coordination structure whenever priorities change, but ensuring that the existing structure continues to adapt as federation matures.

8.3 Funding and sustaining federation

Investing in federation

Over £2 billion is already committed to UK compute infrastructure through the AI Research Resource, National Compute Resources, and the Next National Supercomputing Service. Alongside these national investments, universities, research institutes,



and national facilities continue to invest substantially in local compute, data, and AI capabilities.

Federation does not principally create new infrastructure. Its purpose is to increase the accessibility, interoperability, and collective value of infrastructure that already exists and that of future investments. Federation should be understood as an enabling investment where the return on investment is realised through improved utilisation, reduced duplication, broader participation, and a lower barrier to adoption for new research communities.

Without federation, the benefits of both the national and institutional investments remain concentrated among organisations already able to access or operate advanced digital research infrastructure. Federation provides a mechanism for extending the reach of those investments across a broader range of institutions and research communities.

Activating federation is estimated at around 60 FTE-years in phase 1, 48 in phase 2 and 20 in phase 3 (128 in total) across the three roadmap phases, corresponding to approximately £16 million of effort over the lifetime of the programme using typical UK full-economic-cost assumptions. These are intended to be illustrative of scale instead of a formal costing for a business case. The central observation is that federation activation represents a relatively small additional investment compared with the national and institutional infrastructure base it seeks to connect and coordinate. The proposition offered here is therefore not that federation creates new infrastructure, but that it increases the value realised from infrastructure that has already been funded.

The purpose of this investment is not simply to support existing large national facilities. It is to extend participation to institutions that could not otherwise afford to federate, addressing directly the risk that federation's benefit concentrates among those already equipped to act alone, and turning the sector's existing, currently uncoordinated investment into shared national capability.

Once a pathway is approved, delivery begins with the mechanisms described in Chapter 5. The working groups are convened and take ownership of the shared foundations. The testing environments are established, and begin closing the knowledge gaps that can be settled empirically. The first decision gate then falls at the end of phase 1, at which point the foundations are either in production at the agreed number of sites or the following phase is re-scoped in light of the evidence generated. This ensures that later investment is released against demonstrated progress.

Modalities: sustainable approaches to funding

The roadmap defines the functions required to coordinate federation but not the organisational form through which they must be delivered. Several established models already exist within UKRI's own portfolio and internationally, representing different risk appetites and different trade-offs across funding, operational support and cov-



erage. The examples below are offered as precedents, demonstrating that federation coordination can be implemented through more than one institutional model.

A *programme model*, exemplified by DARE UK, is funded upfront by UKRI across defined phases and requires a single, committed multi-year funding decision. It is co-delivered by two existing, established organisations, HDR UK and ADR UK, who jointly oversee delivery through chartered working groups (26, 27), which lowers start-up risk since existing governance and infrastructure are reused. Coverage is defined by those chartered working groups, with an agreed scope and lifespan; it can expand by chartering further groups, but starts bounded rather than comprehensive.

A *dedicated national body*, exemplified by Canada's Digital Research Alliance, is fully government-funded from inception as standing national infrastructure, with substantial dedicated funding committed at the outset. It is a purpose-built stand-alone organisation with its own governance from the outset (9), which carries a higher initial set-up cost and time since nothing is inherited from an existing body. Its coverage is the broadest of the three from day one: a membership model spanning over 120 institutions, designed for comprehensive national reach. This is not without cost: a government evaluation of the Alliance's own transition to this model found genuine difficulty moving from a decentralised structure to a centralised one while simultaneously upgrading services, with resulting delays and administrative overhead that has taken several years to bring down. A UK equivalent starting from NFCS's existing structure, rather than a comparable decentralised starting point, would not necessarily face the same transition — but the risk is real enough to plan for, not assume away.

An *opportunistic model*, exemplified by the Software Sustainability Institute (now the Research Software Institute), which started with funding from a single research council and grew incrementally over four phases as further councils joined, keeping the upfront commitment low and letting funding follow demonstrated value. It operates as a small, distributed team across several universities (28), which is lean and low-overhead but limited in its capacity for large-scale coordination in its early phases. Its coverage started narrow and broadened over years; it is now funded by all seven UK research councils, but that breadth took fifteen years to reach.

All three models provide mechanisms for extending participation beyond institutions already equipped to act independently, but they differ in how quickly coverage can be achieved and how much investment is required upfront. The programme model resources smaller institutions to participate through funded, chartered activity from the outset. The dedicated national body model gives them a permanent, formal seat in a standing organisation built around their membership, at the cost of a larger initial commitment and set-up time. The opportunistic model lowers the cost of getting started at all, at the cost of coverage that takes years, not months, to broaden.

Each is compatible with the roadmap's working-group structure and shared-



foundations approach.

The roadmap does not prescribe which of these models should be adopted. That decision rests with funders and boards. What the evidence demonstrates is that the proposed working-group structure, shared testing environments, and coordination functions are compatible with several proven approaches to national coordination. The question is not whether such a function is needed, but how it should be funded and governed.

8.4 Final words

The roadmap does not answer every question. It records the uncertainties that remain, identifies the mechanism through which they can be resolved, and proposes pathways that are consistent with the strongest areas of evidence available today. The central conclusion is clear: federation is not a future aspiration but a practical next step. The UK possesses most of the technical, governance, and community foundations required. The challenge is not whether the foundations can be established but rather how they can be coordinated, sustained, and allowed to grow into a coherent national capability. This roadmap offers an evidence-based community-driven route by which that can happen.



Bibliography

[1] Government Office for Science. Large-scale computing: The case for greater UK coordination. Report, with a preface by Sir Patrick Vallance, Government Chief Scientific Adviser, 2021. URL https://assets.publishing.service.gov.uk/media/654a4025e2e16a000d42aaef/UK_Computing_report_-_Final_20.09.21.pdf.

[2] Future of Compute Review panel. Independent review of the future of compute. Independent report, commissioned by the UK Government, 2023. URL https://www.gov.uk/government/publications/future-of-compute-review.

[3] UK Research and Innovation. The UK's research and innovation infrastructure: Opportunities to grow our capability. Report, with forewords by Professor Mark Thomson (STFC) and Sir Mark Walport (UKRI), 2020. URL https://www.ukri.org/wp-content/uploads/2020/10/UKRI-201020-UKinfrastructure-opportunities-to-grow-our-capacity-FINAL.pdf.

[4] Department for Science, Innovation and Technology and UK Research and Innovation. Uk compute roadmap. Policy paper, CP 1352, 2025. URL https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap.

[5] Jisc. Mapping federation journeys for optimising the UK digital research infrastructure. Report, commissioned by the UKRI Digital Research Infrastructure Committee (DRIC), 2024. URL https://www.jisc.ac.uk/reports/mapping-federation-journeys-for-optimising-the-uk-digital-research-infrastructure.

[6] UK Research and Innovation and Department for Science, Innovation and Technology. AI research resource (AIRR). Programme page, 2025. URL https://oecd.ai/en/dashboards/policy-initiatives/ai-research-resource-airr.

[7] UK Research and Innovation. National compute resources (NCR). Programme announcement, 2026. URL https://www.ukri.org/news/new-national-computing-resources-to-open-doors-for-researchers/.

[8] UK Research and Innovation. The UK's next national supercomputer. Programme page, hosted by EPCC, University of Edinburgh, 2026. URL https://www.ed.ac.uk/supercomputing.

[9] Digital Research Alliance of Canada. The digital research alliance of canada. Organisation website, 2026. URL https://www.alliancecan.ca/en/about.

[10] European Commission Joint Research Centre. Open but not powerless: Towards a common understanding of EU digital sovereignty. Policy brief, 2026. URL https://publications.jrc.ec.europa.eu/repository/handle/JRC144908.

[11] EuroHPC Joint Undertaking. Eurohpc federation platform (EFP). Programme page, 2026. URL https://www.eurohpc-ju.europa.eu/first-release-eurohpc-federation-platform-streamline-access-europes-supercomputing-resources-2026-04-15_en.

[12] European Open Science Cloud. The european open science cloud (EOSC). Programme and federation of research infrastructure, 2026. URL https://eosc.eu/eosc-about.

[13] ELIXIR. ELIXIR: The european life-sciences infrastructure for biological information. Research infrastructure, 2013. URL https://elixir-europe.org/.

[14] Worldwide LHC Computing Grid. The worldwide lhc computing grid (wlcg). URL https://wlcg.web.cern.ch/. International research infrastructure.

[15] US National Science Foundation. Advanced cyberinfrastructure coordination ecosystem: Services & support (ACCESS). Programme page, 2022. URL https://access-ci.org/about/.

[16] US Department of Energy, Office of Science. Integrated research infrastructure (IRI). Programme, 2023. URL https://iri.science/about.

[17] Nationale Forschungsdateninfrastruktur (NFDI). National research data infrastructure (nfdi). URL https://www.nfdi.de/en/. German national research data infrastructure.

[18] GÉANT. eduGAIN. Interfederation service, 2026. URL https://edugain.org/.

[19] IRIS. Iris e-infrastructure for research and innovation at stfc. URL https://www.sc.stfc.ac.uk/programmes/iris/. Community-driven federated digital research infrastructure programme.

[20] DiRAC. Dirac high performance computing facility. URL https://dirac.ac.uk/. UK HPC Service Provider for theoretical cosmology, nuclear physics, astrophysics, particle physics, and solar and planetary science.

[21] GridPP Collaboration. Gridpp: Distributed computing for data-intensive research. URL https://www.gridpp.ac.uk/. UK distributed research computing infrastructure.

[22] AARC (Authentication and Authorisation for Research and Collaboration). AARC blueprint architecture 2025 (AARC-BPA-2025). Guideline AARC-G080, 2025. URL https://aarc-community.org/guidelines/aarc-g080/.

[23] National Cyber Security Centre. Cyber assessment framework. Version 4.0, 2025. URL https://www.ncsc.gov.uk/collection/cyber-assessment-framework.



[24] Department for Science, Innovation and Technology. National data library: Progress update. Progress update, 2026. URL https://www.data.gov.uk/.

[25] NetDRIVE. Network for sustainable digital research infrastructure vision and expertise (NetDRIVE). UKRI-funded network project, 2025. URL https://uknetdrive.org/.

[26] DARE UK. DARE UK (data and analytics research environments UK). Programme website, 2026. URL https://dareuk.org.uk/how-we-work/.

[27] Health Data Research UK. Health data research UK (HDR UK). Institute website, 2025. URL https://www.hdruk.ac.uk/.

[28] Software Sustainability Institute. The UK software sustainability institute (SSI). Institute website, 2024. URL https://www.software.ac.uk/about/our-funders.



Tracing Projects Across Recommendations

This appendix maps each umbrella recommendation to the flexible-funded projects and individual project recommendations from which it was consolidated.

Deliver First—see §4.1

B1 Common Data Movement, Cataloguing Access and Storage Standard

Project	Recommendation(s)
ACCoRD	ACCORD-01,02,03
Exploring a Data Centre API for Federated Data Movement	DCAPI-01
Federated Data Access Across the DRI	FACCESS-01, 02, 04, 05
Federated Data Movement	FMOVE-01, 05, 09
Research Data Landscape Survey	RDLs-02, 03, 05
Towards Federated Laboratory Data Management	LABSCI-02, 04, 05, 06
Everything Everywhere All at Once	EEAAOO-08
Enhancing HPC Adoption through User Centred Design	UCDSN-04

D1 Federated Identity and Access Management (IAM)

Project	Recommendation(s)
Federated IAM for Existing Infrastructures	FIAM-01, 02, 03, 05
Exploring Governance for MyAccessID	MYAID-03, 04, 05
Job Orchestration on Heterogenous Systems	CONST-01
Everything Everywhere All at Once	EEAAOO-03
Federated Edge-HPC Architectures for AI Workflows	FEDGE-06
Evaluation of Secure Federated Kubernetes Storage for TREs	FEDK8S-04
Towards Federated Laboratory Data Management	LABSCI-02, 04
Research Data Landscape Survey	RDLS-02
Federated AI Container Registry Feasibility Study	AICON-07

D2 Coordinated Cybersecurity Governance and Operations

Project	Recommendation(s)
DRI Cybersecurity	CYBER-01, 03, 04, 05

D3 Federation Governance and Operating Model

Project	Recommendation(s)
Inclusive Futures	INCL-01, 03, 05
UNITED	UNITED-01, 02, 04
Everything, Everywhere, All at Once	EEAAO-06
Supercomputing and Superpositions	QUANTUM-09
Enhancing HPC Adoption through User Centred Design	UCDSN-04

E1 Single Front Door, Onboarding, Admin and Support

Project	Recommendation(s)
Bridging the Gap - aligning project admin with access to DRI	BGAP-01, 02, 03, 07,
Enhancing HPC Adoption through User Centred Design	UCDSN-01, 07
Barriers for Biology Researchers	BIO-01, 02
Surveying Accessibility for Junior Researchers	JUNIOR-02
Supercomputing and Superpositions	QUANTUM-07
Everything, Everywhere, All At Once	EEAAO-02

E3 Culture, Widening Participation and User Needs

Project	Recommendation(s)
Federation in the Arts and Humanities	A&H-01, 03, 04
Barriers for Biology Researchers	BIO-03, 04
Inclusive Futures	INCL-01
UNITED	UNITED-03

Enablers—see §4.2

A2 Federated Access Mechanisms

Project	Recommendation(s)
Tokenised Authentication for Research Computing Services	TARCS-01, 02, 03
Federation in the Arts and Humanities	A&H-01
Barriers for Biology Researchers	BIO-03

B3 Federated Trusted Research Environments

Project	Recommendation(s)
Evaluation of Secure Federated Kubernetes Storage for TREs	FEDK8S-01, 02, 07, 12

E2 Skills, RSE and Communities of Practice

Project	Recommendation(s)
DRI Cybersecurity	CYBER-04
Everything, Everywhere, All at Once	EEAAO-10
Federated Data Movement	FMOVE-07
Federated Edge-HPC Architectures for AI Workflows	FEDGE-08
ACCoRD	ACCoRD-02
Barriers for Biology Researchers	BIO-02
Surveying Accessibility for Junior Researchers	JUNIOR-01

Deliver Later—see §4.3

A1 Federated Workflows and Job Orchestration

Project	Recommendation(s)
Job Orchestration on Heterogenous Systems	CONST-04, 05
Federated Edge-HPC Architecture for AI Workflows	FEDGE-03, 04, 07
Towards Federated Laboratory Data Management	LABSCI-05, 07
Federation in the Arts and Humanities	A&H-02
Enhancing HPC Adoption Through User Centred Design	UCDSN-03

A3 Industry and Public Sector Access

Project	Recommendation(s)
Job Orchestration on Heterogenous Systems	CONST-03
Enhancing HPC Adoption through User Centred Design	UCDSN-01, 04
Federated Edge-HPC Architecture for AI Workflows	FEDGE-09

C1 Federated Software and Container Supply Chain

Project	Recommendation(s)
Federated AI Container Registry Feasibility Study	AICON-01, 02, 05, 06, 11, 12
Enhancing HPC Adoption through User Centred Design	UCDSN-03
DRI Cybersecurity	CYBER-02
Federated Edge-HPC Architecture for AI Workflows	FEDGE-02



C2a AI Research-Support Services across the Federation

Project	Recommendation(s)
Federated AI Helpdesk	AIHELP-01, 02, 03
Federation in the Arts and Humanities	A&H-01

C2b Responsible AI Governance

Project	Recommendation(s)
Federated AI Helpdesk	AIHELP-04
Federation in the Arts and Humanities	A&H-02
DRI Cybersecurity	CYBER-02

F1 Fair, Transparent Resource Allocation and Accounting

Project	Recommendation(s)
FAIR Compute	FAIRC-01, 02, 06
Enhancing HPC Adoption through User Centred Design	UCDSN-07

Horizon—see §4.4

C3 Quantum–HPC Integration Pathway

Project	Recommendation(s)
Supercomputing and Superpositions	QUANTUM-01, 02, 03, 04, 08, 09

F2 Carbon and Resource Aware Federation

Project	Recommendation(s)
Net Zero and Circular Economy Federation	NETZERO-01, 03, 05



Flexible-funded Projects and their Outputs

Table 1: Phase 1 Projects and Outputs

Project (Short Code)	Output
ACCoRD: A Community For Contract Regulation For Data For The Project Name (ACCORD)	https://doi.org/10.5522/04/33103166
Bridging the Gap: Aligning Project Administration with Access to Digital Research Infrastructure (BGAP)	https://doi.org/10.5281/zenodo.21706698
Enhancing HPC Adoption Through User-Centred Design: A Roadmap for Inclusive Innovation in Environment, Health, and the Built Environment (UCDSN)	https://doi.org/10.5281/zenodo.21629487
Exploring the governance requirements for enabling UK DRIs to adopt MyAccessID (MYAID)	https://doi.org/10.5281/zenodo.21718790
Exploring the requirements and technologies for a data centre API for federated data movement (DCAPI)	https://doi.org/10.5281/zenodo.21870345
Federated AI Application Container Platform/Registry Feasibility Study (AICON)	https://doi.org/10.5281/zenodo.21701682
Federated Data Movement (FMOVE)	https://doi.org/10.5281/zenodo.21710495
Federated Edge–HPC Architectures for AI Workflows in Privacy-Sensitive and Real-Time Domains (FEDGE)	https://doi.org/10.5281/zenodo.21940454
Federated IAM for existing infrastructures (FIAM)	https://doi.org/10.7488/era/7458
Federation of Compute and Infrastructures in the Arts and Humanities (A&H)	https://doi.org/10.5281/zenodo.21706879
Inclusive Futures: User stories and Mapping Pathways for National Federated Compute Services (INCL)	https://doi.org/10.5281/zenodo.21673355
UKRI Research Data Landscape Survey (RDLS)	https://doi.org/10.5281/zenodo.21724665
UNITED (User Needs Informed Technologies for Environmental Data): A Framework for Federated Computing Roadmaps (UNITED)	https://doi.org/10.5281/zenodo.21642608

Table 2: Phase 2 Projects and Outputs

Project (Short Code)	Output(s)
A federated bespoke AI-assisted helpdesk for DRI facilities (AIHELP)	https://doi.org/10.5281/zenodo.21710690
DRI Federation Cybersecurity Roadmap development (CYBER)	https://doi.org/10.5281/zenodo.21933720
Evaluation of Secure Federated Kubernetes Storage for Trusted Research Environments (FEDK8S)	https://doi.org/10.5281/zenodo.22831674
FAIR-Compute: A Roadmap for Fair and Efficient Allocation of Federated Digital Research Infrastructure (FAIRC)	https://doi.org/10.48550/arXiv.2607.28290
Federated data access across the DRI (FACCESS)	https://doi.org/10.5281/zenodo.21708310
Federating Everything, Everywhere, All at once – exploring the potential for inclusion of Tier3 services in the UKRI federated DRI (EEAAO)	https://doi.org/10.25392/leicester.data.33123347
Identifying Barriers for Biology Researchers Using Federated HPC Services (BIO)	https://doi.org/10.5281/zenodo.21680185
Job Orchestration using Constellation on Heterogenous HPC resources (CONST)	https://doi.org/10.25392/leicester.data.33123338
Net-Zero and Circular Economy Federation: Evidence-based Policy Roadmap for Carbon-Aware Compute in the Built Environment (NETZERO)	https://doi.org/10.5281/zenodo.21709641
Supercomputers and Superpositions: Making Quantum Accelerators Accessible Within HPC Frameworks (QUANTUM)	https://doi.org/10.5281/zenodo.21776313
Surveying accessibility of federated compute for junior researchers (JUNIOR)	https://zenodo.org/records/21918401
Tokenised Authentication for Research Computing Services (TARCS)	https://doi.org/10.7488/era/7519
Towards a common data infrastructure for laboratory science (LABSCI)	https://doi.org/10.6084/m9.figshare.33349557



Areas of contention and compromise

The B1/B2 merger The July 2026 consultation judged that **B2** (storage, cataloguing and curation) belonged inside **B1**, since separating movement from storage and cataloguing was artificial. The umbrellas were merged without renumbering, preserving traceability to the underlying project evidence. One residual: the merge left cataloguing without a clearly named home within **B1**, which is why it also appears as a standalone funding proposition in the data roadmap (§7.5) and as a recorded knowledge gap (**sto-g4**).

The governance band: early but non-gating Consultation feedback pulled in two directions on governance: one view held it should start early, given how much depends on it; another held it should not gate everything else, risking delivery stalled behind a slow-to-agree framework. Both are reflected in **D1**, **D2** and **D3**: placed in the deliver-first band as an early, enabling group, but not treated as a single gate other work must wait behind. This is a considered compromise rather than a resolution, and may resurface as the roadmap moves from strawman to delivery.

Selecting among external standards and efforts The roadmaps repeatedly reference external reference points — eduGAIN, EuroHPC (11), EOSC (12), ELIXIR (13), NSF ACCESS (15), DOE's IRI (16), the UK Compute Roadmap (4) and the National Data Library (24) — treating interoperation with them as a benefit in its own right. The document does not state a basis for choosing between them where standards, timescales or priorities diverge, or which to prioritise when resources don't allow all at once. In practice this selection is expected to fall to the federation working groups and be tested empirically in the shared testing environments, but that is an implicit expectation rather than a stated policy, not yet put to communities or boards as a question in its own right. A related, more immediate tension sits closer to home. The AI Research Resource (6) and National Compute Resources (7) are UK national compute investments with access and governance arrangements already in motion rather than proposed. NFCS's federation model, particularly the participation levels in §2.3, has not been tested against either programme's existing access model, and neither was designed with these levels in mind. Whether alignment is straightforward, needs negotiation, or reveals a genuine incompatibility isn't yet known, more pressing than the international comparisons, since both are live now.

Glossary

Assumption A stated planning basis (timing, cost basis, or scope and method) underlying the roadmap options, numbered `asm-1` to `asm-9` and cited from the option, appraisal, and phase tables so it can be reviewed and challenged independently rather than treated as settled (Table 5.3).

Conformance The entry profile a service meets in order to participate in federation — for example, to the data-movement and access standard or the shared cybersecurity baseline — verified rather than assumed.

Federated Widely used throughout this document and in national strategy, but without an agreed operational definition against which a service can be assessed. This is recorded as a knowledge gap (**fed-g1**); §8 sets out why settling it is both answerable and inexpensive to fund early. What this document *can* define — the levels at which a service can participate — is set out in §2.3.

Federation A framework through which independently governed services and infrastructures operate as part of a coherent national ecosystem, enabling common access, interoperability, coordination and shared standards while participants retain control over their own infrastructure.

Federation coordination function The continuing role of overseeing the federation, maintaining cross-cutting standards and supporting adoption as new communities and investments engage, distinct from the delivery of any single roadmap capability.

Knowledge gap A question the programme's evidence could not settle, recorded with an identifier so it can be tracked and, where appropriate, funded separately. Gaps are numbered `fed-g1` to `fed-g7` for compute federation and `sto-g1` to `sto-g5` for data services (Table 3.3).

Level 0: Independent A service operates on its own community's identity, access and data arrangements, with no shared standards adopted.

Level 1: Interoperable A service adopts common identity, access and data-interoperability standards, making it discoverable and usable by other federated services without further operational change.

Level 2: Integrated A service exposes its own resources, data and operations through federation mechanisms directly.

Pillar One of the programme's three organising dimensions — Community, Governance, and Technology — used to check that no aspect of federation was overlooked when evidence was gathered.

Priority band The four groupings into which umbrella recommendations are sorted on the priority × urgency map (Figure 3.2): *deliver-first* (must start now and everything else depends on it), *enablers* (can proceed alongside the foundations, offering early value), *deliver-later* (high priority but dependent on the foundations

being in place first), and *horizon* (strategically important but dependent on longer-term technology or maturity).

Shared foundation A capability that both the compute and data roadmaps depend on and that is delivered once for both, because neither roadmap can deliver it independently (§5.1). Not the same set as the deliver-first band: three recommendations (D2, E1, E3) are deliver-first without being shared, and three (A2, B3, F1) are shared without being deliver-first.

Testing environment A shared facility used to settle knowledge gaps empirically, verify conformance to the shared foundations, and de-risk capabilities before they are scaled — not a production national service.

Theme One of six groupings (A–F) that structure the 17 umbrella recommendations for navigability, reflecting that most projects span more than one pillar.

Umbrella recommendation One of the 17 consolidated groupings of the programme's 259 individual project recommendations, each with full traceability to the underlying evidence (Table 3.2).

Working group One of four cross-cutting groups, each aligned to a shared foundation, responsible for the ongoing standards, governance and conformance criteria that let independently governed services interoperate.

