SITRA

WORKING PAPER

SEPTEMBER 2026

THE NEXT RIGHT QUESTIONS ON THE PATH TO THE AGENTIC STATE

Functions as the shared level for public-sector AI decisions, learning and renewal

Jonna Heliskoski
Programme Director
Sitra

Tommi Vihervaara
Senior Lead
Sitra

Jari Ängeslevä
Senior Lead
Sitra

Antti Harjunpää
Senior Lead
Sitra

Kirsi Hanhisalo
Senior Lead
Sitra

Many public-sector organisations are experimenting with AI, yet more pilots do not add up to public sector renewal: similar work is still being redesigned separately in different organisations. The next question is not which task to automate or which AI tools to adopt, but at what level public work should be understood, compared and redesigned. This working paper proposes the function – a recurring activity with a recognisable result, such as receiving an application or issuing a license – as that shared level.

The working paper is published in tandem with First Moves: A Finland-Specific Assessment of Agentic AI Potential in Government – a study by the Agentic State foundation and Sitra – which maps where recurring public-sector work may be technically and legally suitable for agentic AI. This paper asks how that emerging map can be turned into decisions about what work to redesign, test and scale first.

Sitra's working paper, September 2026

The Next Right Questions on the Path to the Agentic State – Functions as the shared level for public-sector AI decisions, learning and renewal

Authors:
Jonna Heliskoski, Programme Director, Sitra
Antti Harjunpää, Senior Lead, Sitra
Tommi Vihervaara, Senior Lead, Sitra
Kirsi Hanhisalo, Senior Lead, Sitra
Jari Ängeslevä, Senior Lead, Sitra

Layout: Grano Oy

ISBN 978-952-347-483-3 (PDF) www.sitra.fi
ISSN 2737-1042 (online publication) www.sitra.fi

Sitra working papers provide multidisciplinary information about developments affecting societal change. Working papers are part of Sitra's future-oriented work conducted by means of forecasting, research, projects, experiments and education.

Contents

Summary 5
Tiivistelmä 7
Sammanfattning 8
1. At what level should AI-enabled renewal in the public sector begin? 9
2. What exactly is a function? 11
2.1 Five levels for understanding the same work 12
2.2 Where can an AI agent act? 13
2.3 Functions and life events describe different views of the same system 13
3. Which functions should be prioritised? 14
3.1 A shared language makes comparison possible 15
4. What must change before AI can improve productivity? 16
4.1 Data must be understandable, not merely available 16
4.2 Build around durable work, not temporary systems 17
4.3 Human responsibility must be designed into the workflow 17
4.4 Shared capability can also scale failure 18
4.5 Five gates before technology becomes a productivity investment 18
5. How can the Agentic State learn across organisations? 20
5.1 Four leadership responsibilities 20
5.2 Learning needs to travel through three levels 21
5.3 What should be learned from the first demonstrator? 21
6. How can shared functions create reusable capability and a GovTech market? 23
6.1 Different services can share part of the same work 23
6.2 From shared work to shared capability 23
6.3 From shared capability to a GovTech market 24
6.4 Shared infrastructure beneath reusable capabilities 25
6.5 Existing examples show parts of the logic 25
6.6 What would this mean for application intake? 25
7. How should shared functions be governed across organisational boundaries? 26
7.1 Start with a coordination mandate, not new institutional ownership 26
7.2 Four practical problems must still be solved 27
7.3 What can be done within the next 18 months? 27

7.4 Governance must follow the level of the decision 28
7.5 Towards 2030: Some questions will require structural decisions 28
8. What should leaders do next? 30
8.1 Decide now 30
8.2 Prepare next 30
8.3 Do not decide yet 30
8.4 Decide towards 2030 30
8.5 The practical starting point 31
Essential terms 32
References 33

Summary

Many public-sector organisations are experimenting with artificial intelligence, yet more pilots do not automatically add up to public-sector renewal. Similar work is still redesigned and procured separately across organisations, cases stall at organisational boundaries, and people are repeatedly asked for information the public sector already holds. The next question is therefore not simply which AI tools government should adopt, but *at what level public work should be understood, compared and redesigned.* This working paper proposes the **function** as a shared level for doing so.

A function is a recurring public activity with a clear result, such as receiving an application, assessing eligibility, issuing a licence, making a payment, or providing guidance. The same function may cross organisational boundaries or recur in similar form across many public bodies.

Agentic AI changes what is possible and this builds the case for function-based thinking. AI agents can increasingly carry out bounded, multi-step work across systems and organisational boundaries. That capability creates an opportunity to redesign recurring public work as a whole rather than adding new technology to existing organisational processes.

Function-based thinking does not remove organisations, legal mandates, or accountability. It makes repeated work visible across them. It allows leaders to distinguish what can be shared from what should remain specific, identify where reusable capability could create value, and coordinate dependencies that no organisation can solve efficiently alone.

The function links strategic priorities to operational work. It sits below public purpose and portfolios and above workflows and tasks. That role makes it the level at which leaders can compare recurring work and decide where to focus attention, while the detailed human–agent boundary is defined within the workflow. Statutory responsibility remains with the responsible public body.

Function-based thinking also changes how AI opportunities should be prioritised. Technical feasibility alone is not enough. Leaders need to consider value, scale, feasibility and deliverability together: Is the function worth changing? Is the opportunity large enough? Where can AI safely contribute? And can the participating organisations actually deliver the change? Not every technically feasible opportunity is worth pursuing.

The broader opportunity goes beyond automation. Shared definitions, interfaces, safeguards and evidence can allow several public bodies to reuse parts of the same capability while retaining justified variation. Evidence from human–agent work can also travel across organisations, allowing improvements in one implementation to improve the same recurring function elsewhere. In this way, learning from one organisation can create value beyond that organisation.

Before deciding where to invest, leaders first need to understand where agentic AI may be suitable. Published with The Agentic State foundation alongside this working paper, *First Moves: A Finland-Specific Assessment of Agentic AI Potential in Government* maps where recurring public work may be technically and legally suitable for agentic AI. That map is a starting point, not an investment agenda: it does not show actual workflows, costs, benefits, service outcomes or organisational readiness. This working paper asks the next right questions: *How should government turn that emerging map of potential into decisions about what to redesign, test and scale?*

This paper proposes that the next 18 months should be treated as a bounded learning period, not as a race to build a national platform. The objective is to learn what can be

reused, governed and scaled, not to prove a predetermined solution. To do this, Finland should select one recurring function with limited legal discretion and a clear boundary for human judgement. Participating organisations should establish the baseline, redesign the workflow and data before procuring technology, and test whether a reusable capability can work across several public bodies.

Application intake, including initial processing, is a credible candidate for such a function because it recurs across the public sector, contains measurable repetition and delay, and can be separated from the final legal decision.

By 2030, Finland will need to answer harder questions about permanent stewardship, funding, standards and accountability. For now, the practical task is simpler: Find one important function that several public bodies already perform, make the shared work visible, and decide what those public bodies should redesign, learn or build together.

How to read this working paper

This working paper can be read in three parts:

- **The case for functions**: Sections 1 and 2 explain why functions are useful for understanding recurring public work and define the key concepts.
- **Choosing the first test**: Sections 3 and 4 set out how to choose and design the first demonstrator.
- **Learning, reuse and governance**: Sections 5 to 8 explain how learning, reusable capability, governance and next decisions should develop from that first test.

Tiivistelmä

Useissa julkisen sektorin organisaatioissa kokeillaan tekoälyn käyttöä, mutta pilottihankkeiden määrän kasvu ei automaattisesti tarkoita julkisen sektorin uudistumisen etenemistä. Samanlaisia uudistuksia suunnitellaan ja hankitaan erikseen eri organisaatioissa, asiat seisahtuvat organisaatiorajoille ja asiakkaina olevilta kansalaisilta pyydetään toistuvasti tietoja, jotka julkisella sektorilla on jo hallussaan. Tekoälyn käyttöönotossa, olennainen seuraava kysymys julkiselle sektorille ei siksi ole pelkästään määrittää tekoälytyökalut, jotka valtionhallinnossa tulisi ottaa käyttöön, vaan ensin *määrittää tarkastelutaso, jolla julkisen sektorin työtä tulisi tarkastella, vertailla ja suunnitella uudelleen.* Tässä työpaperissa ehdotamme toimintoa (function) täksi tarkastelutasoksi.

Toiminto tarkoittaa toistuvaa julkisen sektorin toimintaa, jolla on selkeä tulos, kuten hakemuksen vastaanottaminen, kelpoisuuden arviointi, luvan myöntäminen, maksun suorittaminen tai neuvonnan antaminen. Sama toiminnon toteuttaminen voi ylittää organisaatiorajoja tai se voi esiintyä samanlaisessa muodossa useissa eri julkisen sektorin organisaatioissa.

Agenttinen tekoäly muuttaa sitä, mikä on mahdollista ja tämä tukee toimintoperustaista ajattelutapaa. Tekoälyagentit pystyvät kasvavassa määrin suorittamaan rajattuja, monivaiheisia tehtäviä ylittäen järjestelmien ja organisaatioiden rajoja. Tämä kyky tarkoittaa sitä, että on mielekästä suunnitella eri organisaatioissa toistuva työ ja niiden kulku ensin kokonaan uudelleen, sen sijaan, että teknologiaa lisättäisiin nykyisiin organisaatiokohtaisiin prosesseihin.

Toimintoperustainen ajattelutapa ei poista organisaatioita, lakisääteisiä toimivaltuuksia tai vastuuvelvollisuutta. Se tuo näkyviin samanlaisena toistuvan työn julkishallinnossa. Se antaa päätöksentekijöille mahdollisuuden erottaa, mikä julkisen sektori työssä on yhteistä ja mikä tulisi pitää organisaatiokohtaisena.

Toimintoperustainen lähestymistapa auttaa tunnistamaan missä uudelleenkäytettävä kyvykkyys voi luoda arvoa, sekä auttaa koordinoimaan riippuvuuksia, joita mikään yksittäinen organisaatio ei voi ratkaista tehokkaasti yksin.

Tarkastelutasona toiminto yhdistää strategiset painotukset ja operatiivisen työn. Toiminto on kapeampi taso kuin organisaation yleinen tarkoitus ja vastuualueet, mutta laveampi kuin työnkulut ja tehtävät. Toiminnon tasolla voidaan vertailla samankaltaisena toistuvaa työtä ja päättää uudistamisen painopistealueista. Ihmisen ja agentin työn välinen raja puolestaan määritellään yksityiskohtaisesti työnkulun tasolla. Lakisääteinen vastuu säilyy asianomaisella julkisella organisaatiolla.

Toimintopohjainen ajattelu muuttaa myös sitä, miten tekoälyn tarjoamat mahdollisuudet tulisi priorisoida. Pelkkä tekninen toteutettavuus ei riitä. Johtajien on otettava huomioon arvo, laajuus, toteutettavuus ja toimitettavuus kokonaisuutena: Kannattaako toimintoa uudistaa? Onko hyötypotentiaali riittävän suuri? Missä tekoälyä voidaan turvallisesti hyödyntää? Pystyvätkö mukana olevat organisaatiot saamaan muutoksen aikaan? Kaikki teknisesti toteutettavissa olevat mahdollisuudet eivät ole toteuttamisen arvoisia.

Lähestymistapa tarjoaa tiettyjen tehtävien automatisointia laajempia mahdollisuuksia. Yhteiset määritelmät, rajapinnat, suojatoimet ja kertyvä näyttö voivat antaa useille julkisen sektorin organisaatioille mahdollisuuden hyödyntää osittain samaa kyvykkyyttä, säilyttämällä kuitenkin perustellut organisaatiokohtaiset erityisyydet. Ihmisten ja agenttien yhteistyöstä saadut kokemukset voivat myös levitä organisaatiosta toiseen, jolloin yhdessä toteutuksessa tehdyt parannukset voivat tehostaa samanlaisena toistuvaa toimintoa myös muualla. Tällä tavoin organisaatiossa saadut oppimiskokemukset voivat tuottaa arvoa myös sen ulkopuolella.

Ennen investointipäätöksiä johdon on ymmärrettävä, mihin agenttipohjainen tekoäly soveltuu. Tämän työpaperin lisäksi julkaisemme yhdessä Agentic State foundationin kanssa selvityksen *First Moves: A Finland-Specific Assessment of Agentic AI Potential in Government*, jossa luodaan yleiskuva samanlaisena toistuvasta julkishallinnon työstä, jossa agenttisen tekoälyn käyttö voi olla teknisesti ja oikeudellisesti tarkoituksenmukaista. Yleiskuva toimii lähtökohtana, ei investointisuunnitelmana: siinä ei kuvata todellisia työnkulkuja, kustannuksia, hyötyjä, palvelujen tuloksia tai organisaatioiden valmiutta. Tässä työpaperissa esitetään seuraavat olennaiset kysymykset: *Miten valtionhallinnon tulisi hyödyntää tätä tilannekuvaa tehdessään päätökset siitä, mitä tulisi uudistaa, kokeilla ja skaalata?*

Tässä työpaperissa ehdotetaan, että seuraavat 18 kuukautta tulisi nähdä oppimisjaksona, ei kilpajuoksuna kansallisen tekoälyalustan rakentamiseksi. Tavoitteena tulisi olla selvittää, mitä samoja palveluita ja niitä tukevia ratkaisuja voidaan käyttää eri organisaatioissa. Päästäkseen alkuun Suomen tulisi valita yksi samanlaisena toistuva toiminto, jossa lainsäädännöllinen harkintavalta on rajallinen ja ihmisen tekemän arvion rajat ovat selkeästi määriteltyjä. Osallistuvien organisaatioiden tulisi määrittää lähtötilanne, suunnitella työnkulku ja tiedot uudelleen ennen teknologian hankkimista sekä kokeilla, toimiiko uudelleen käytettävä kyvykkyys useissa eri julkisissa organisaatioissa.

Hakemusten vastaanotto, mukaan lukien niiden alustava käsittely, on hyvä vaihtoehto tällaiseksi toiminnoksi, koska se toistuu samanlaisena koko julkisella sektorilla, sisältää mitattavissa olevia toistuvia vaiheita ja viivästyksiä. Se voidaan myös erottaa lopullisesta oikeudellisesti sitovasta päätöksestä.

Vuoteen 2030 mennessä Suomen on vastattava entistä vaikeampiin kysymyksiin, jotka koskevat pysyvää johtamista, rahoitusta, standardeja ja vastuuvelvollisuutta. Käytännön tehtävä on kuitenkin tässä vaiheessa yksinkertainen: Määritä yksi tärkeä toiminto, jota useat julkiset organisaatio jo hoitavat, havainnollista toistuva työ ja päätä, mitä julkisten toimijoiden tulisi uudistaa, oppia tai kehittää yhdessä.

Sammanfattning

Många organisationer inom den offentliga sektorn testar sig fram inom artificiell intelligens, men fler pilotprojekt bidrar inte automatiskt till en förnyelse av den offentliga sektorn. Likartat arbete omformas och upphandlas fortfarande separat inom olika organisationer, ärenden fastnar vid organisationsgränserna och invånarna ombeds upprepade gånger att lämna uppgifter som den offentliga sektorn redan förfogar över. Nästa fråga är därför inte bara vilka AI-verktyg som myndigheterna bör införa, utan *på vilken nivå det offentliga arbetet bör förstås, jämföras och omformas*. I detta arbetsdokument föreslås att en funktion bör fungera som en gemensam utgångspunkt för detta.

En funktion är en återkommande offentlig uppgift med ett tydligt resultat, till exempel att ta emot en ansökan, bedöma behörighet, utfärda ett tillstånd, göra en utbetalning eller erbjuda rådgivning. Samma funktion kan sträcka sig över organisationsgränser eller återkomma i liknande form inom många offentliga organ.

Agentbaserad AI förändrar vad som är möjligt, och detta bygger grunden för funktionsbaserat tänkande. AI-agenter kan i allt högre grad utföra avgränsade arbeten i flera steg som sträcker sig över system- och organisationsgränser. Denna förmåga skapar en möjlighet att omforma återkommande offentliga uppgifter i sin helhet, i stället för att bara utöka befintliga organisatoriska processer med ny teknik.

Ett funktionsbaserat tänkande innebär inte att organisationer, lagstadgade mandat eller ansvarsskyldighet försvinner. Det synliggör återkommande arbetsuppgifter genom alla dessa. Det gör det möjligt för ledare att skilja mellan vad som kan delas och vad som bör förbli specifikt, identifiera var återanvändbar kapacitet kan skapa värde samt samordna beroendeförhållanden som ingen organisation kan lösa effektivt på egen hand.

Funktionen kopplar samman strategiska prioriteringar med det operativa arbetet. Den befinner sig under offentliga syften och portföljer, men ovanför arbetsflöden och uppgifter. Denna roll utgör den nivå, där ledare kan jämföra återkommande arbetsuppgifter och avgöra vad de ska fokusera på, samtidigt som den exakta gränsen mellan människa och agent fastställs inom arbetsflödet. Det lagstadgade ansvaret ligger fortfarande hos den ansvariga offentliga myndigheten.

Funktionsbaserat tänkande förändrar också hur möjligheterna inom AI bör prioriteras. Teknisk genomförbarhet räcker inte som enda krav. Ledare måste beakta värde, omfattning, genomförbarhet och leveransförmåga som en helhet: Är det värt att ändra funktionen? Är möjligheten tillräckligt stor? I vilka sammanhang kan AI bidra på ett säkert sätt? Kan de deltagande organisationerna verkligen åstadkomma förändringen? Inte alla tekniskt genomförbara möjligheter är värda att satsa på.

De större möjligheterna sträcker sig bortom automatisering. Gemensamma definitioner, gränssnitt, säkerhetsåtgärder och underlag kan göra det möjligt för flera offentliga organ att återanvända delar av samma kapacitet samtidigt som berättigade skillnader bevaras. Erfarenheter från samarbete mellan människor och agenter kan också delas mellan organisationer, vilket gör att förbättringar av en viss tillämpning kan leda till förbättringar av samma återkommande funktion på andra håll. På så sätt kan lärdomar från en organisation skapa värde även utanför den organisationen.

Innan en organisation bestämmer sig för var man ska satsa måste ledningen först förstå i vilka sammanhang agentbaserad AI kan vara lämplig. Publikationen *First Moves: A Finland-Specific Assessment of Agentic AI Potential in Government* gavs ut i samarbete med stiftelsen The Agentic State i anslutning till detta arbetsdokument och kartlägger områden där

återkommande offentliga uppgifter kan vara tekniskt och juridiskt lämpliga för agentbaserad AI. Kartläggningen är en utgångspunkt, inte en investeringsplan: den visar inte faktiska arbetsflöden, kostnader, fördelar, tjänsteresultat eller organisationers beredskap. I detta arbetsdokument ställs följande relevanta frågor: *Hur bör regeringen omsätta den framväxande potentialen i beslut om vad som ska omformas, testas och skalas upp?*

I detta arbetsdokument föreslås att de kommande 18 månaderna bör betraktas som en avgränsad inlärningsperiod, inte som en kapplöpning för att bygga en nationell plattform. Målet är att ta reda på vad som kan återanvändas, styras och skalas upp – inte att bevisa att en förutbestämd lösning fungerar. För att uppnå detta bör Finland välja en återkommande funktion med begränsad skönsmässig bedömning och med tydliga gränser för mänskligt omdöme. De deltagande organisationerna bör fastställa utgångsläget, omforma arbetsflödet och data innan de anskaffar teknik, samt testa om återanvändbar kapacitet kan fungera inom flera offentliga organ.

Handläggningen av ansökningar, inklusive den inledande behandlingen, är en lämplig kandidat för en sådan funktion eftersom den återkommer inom hela den offentliga sektorn, innehåller mätbara återkommande moment och fördröjningar samt kan skiljas från det slutliga rättsliga beslutet.

Senast 2030 kommer Finland att behöva ta ställning till svårare frågor om långsiktig förvaltning, finansiering, standarder och ansvarsskyldighet. För tillfället är den praktiska uppgiften enklare: att hitta en viktig funktion som flera offentliga organ redan utför, synliggöra det gemensamma arbetet och bestämma vad dessa offentliga organ bör omforma, lära sig eller bygga upp tillsammans.

1. At what level should AI-enabled renewal in the public sector begin?

Public sector AI decisions cannot be made at one level alone. Governments need to decide why the work matters, where to focus investment, what recurring work could be improved across organisations, and how individual workflows and tasks should change.

Today, the function level is largely missing from the discussion. Most AI initiatives are still framed through organisations, existing services, systems, and procurements. This framing makes local improvements possible, but it also hides repeated work. A ministry, a wellbeing services county, and a municipality may all be improving application intake without recognising that they are addressing parts of the same problem and may need similar data, rules, interfaces and safeguards.

Agentic AI gives leaders a reason to challenge organisation-first framing. Unlike tools designed mainly to support a single task or user, AI agents can increasingly carry out several bounded steps towards a goal. AI agents gather information, check explicit rules, use approved systems, prepare an output and escalate an exception. This means that the relevant unit of redesign may no longer be one task, one system or even one organisation. It may be a recurring sequence of public work that several organisations perform in similar form.

This working paper proposes the *function* as the shared level for understanding and redesigning public-sector work. A function is a recurring public-sector activity, such as receiving an application, assessing eligibility, issuing a licence, making a payment or providing guidance. It may cross organisational boundaries or recur in similar form across several organisations.

Looking at work through functions allows leaders to compare similar activity, identify where reuse may create value and distinguish what can be shared from what should remain specific. Adding the function level does not mean moving every AI decision there, and it does not weaken organisational or legal responsibility. Ministers and senior steering bodies still set the purpose and priorities of public organisations, which retain their budgets, statutory mandates and accountability. Workflow owners and professionals determine how work should operate in practice and where human judgement is required.

The function fills a different gap. It creates a shared way of describing recurring work and asking: *Are we performing essentially the same recurring work? What should remain different? What could we redesign, learn, or build together?*

This way of describing recurring work is particularly relevant to the idea of *the Agentic State:* public administration in which people and AI agents can complete multi-step work across organisational boundaries under clear human authority, shared rules and public-purpose accountability (Ilves et al., 2025).
In this working paper, an AI agent refers to software that can pursue a goal through several bounded steps and escalate when it reaches an exception or a boundary it must not cross. It is not a public official or a legal decision-maker.

The broader report, *First Moves: A Finland-Specific Assessment of Agentic AI Potential in Government*, was prepared jointly by The Agentic State foundation and Sitra. Published in tandem with this working paper, it maps recurring public activities across central government, the wellbeing services counties and the regional-capital municipalities, and scores each function on six dimensions of agentic readiness. It provides an empirical starting point, but not an investment agenda: actual workflows, costs, value, service outcomes, and organisational readiness must still be assessed before priorities are set.

The next 18 months should therefore test a practical proposition: Does looking at recurring public work through functions lead to better decisions about where to redesign, experiment and reuse?

Towards 2030, the question becomes structural: Which decisions should be coordinated at function level, and which must remain with portfolios, organisations, workflows, and professionals?

Next right questions

Short term – next 18 months	**Long term – towards 2030**
Which activity do we perform that several other public bodies also perform, and what better decisions might become possible if we treated it as the same function?	Which decisions should be coordinated at function level, and which must remain with portfolios, organisations, workflows or professionals?

2. What exactly is a function?

A **function** is a recurring public-sector activity that delivers a recognisable result. Examples include receiving an application, checking eligibility, issuing a decision, notifying a person, or making a payment. The same function often recurs in similar form across many public bodies and in some cases the work needed to complete it already crosses organisational boundaries. (World Economic Forum, Capgemini & Global Government Technology Centre Berlin, 2026). A function-based approach keeps recurring work visible even when the organisation, service, or technology around it changes.

Functions give leaders a practical level for understanding and comparing public-sector work. Organisation-level descriptions are often too specific to show where similar work could be reused across organisations. Broad policy areas such as health, education, or social security are too general to redesign. Individual tasks, in turn, are too narrow to show where a larger productivity opportunity may lie.

The function sits between these levels. It is specific enough to connect to real work, but broad enough to reveal where similar work is being performed across organisations.

Four distinctions are important:

- **A function is not an organisation.** A public organisation performs many different functions. Kela, a ministry or a wellbeing services county is therefore not a function.
- **A function is not a policy area.** Health and social care is a policy area; receiving and checking an application is a function within it.
- **A function is not a task.** Reading an address from a document is one task within a workflow that contributes to a function.
- **A function is not a legal mandate.** Recognising similar work as the same function does not change who holds the statutory authority or responsibility to make a decision.

That last distinction is particularly important. Function-based thinking does not transfer authority from one organisation to another. It makes similar work visible so that organisations can compare it, measure it and decide deliberately what could be shared or reused, and what must remain different.

The important shift is not to remove organisations, but to make repeated work visible across them. Organisation-first thinking starts with each public body and can reproduce separate processes, systems and procurements. Function-based thinking asks where the same recurring work appears across organisations, not to make every service identical, but to reveal common parts of the work that may otherwise remain hidden.

Figure 1. Organisation-first thinking duplicates the same work in every public body; function-first thinking makes that shared work visible.

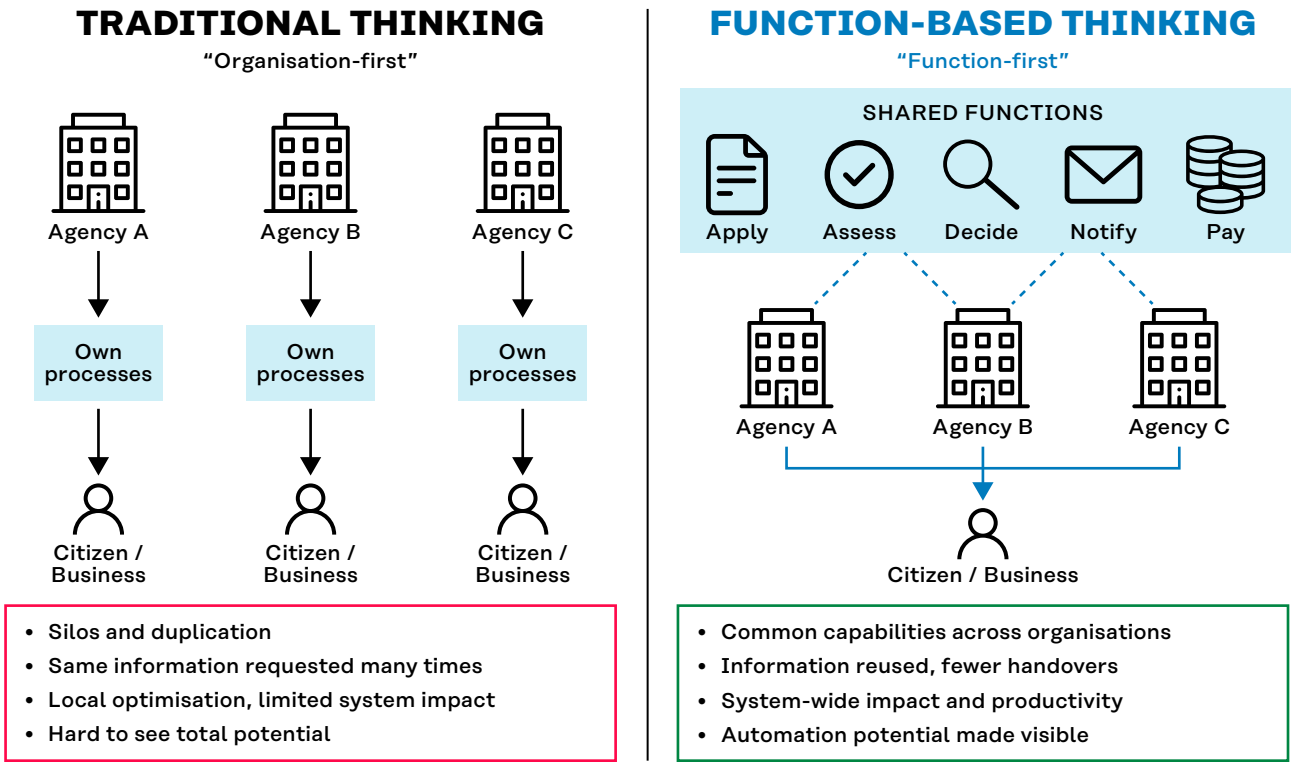

2.1 Five levels for understanding the same work

A function becomes most useful when it is connected both to the public purpose above it and the concrete work below it. To show these connections, the model uses five levels:

Table 1. Five levels for understanding public-sector work through the application-intake example.

Level	Plain-language question	Application-intake example
+2 Public purpose	Why does this work exist?	A person receives the public service or entitlement for which they qualify.
+1 Portfolio	Which related functions should be governed and invested in together?	An access-to-public-services portfolio that groups intake, assessment, decision, notification and payment functions.
0 Function	What recurring result is needed?	Receive an application and prepare it for the responsible decision-maker.
−1 Operational workflow	How does the result get produced?	Receive, verify identity, check completeness, request missing information and route the case.
−2 Task	What specific action is performed?	Read an address, validate a field or draft a request for an attachment.

These levels answer different questions. At the top, public purpose (+2) and portfolios (+1) help political and administrative leaders decide what to prioritise and where to focus attention and investment.

At function level (0), organisations can compare recurring work and ask what could be improved or reused together.

At workflow level (-1), professionals and service owners decide how the work should operate in practice in a specific context, where human judgement is needed and where an AI agent could contribute.

At task level (-2), specific actions can be automated, supported or redesigned.

The model does not reassign legal authority, budgets or accountability, and these remain with the organisations and public authorities that hold them.

2.2 Where can an AI agent act?

The example of application intake makes the distinction concrete. An AI agent might check whether required information is present, retrieve authorised information from an approved source, identify inconsistencies or draft a request for genuinely missing evidence. But that does not mean the entire function has been delegated to an agent.

The boundary between agentic work and human responsibility must be defined within the workflow itself. Where a decision affects a person's rights, the responsible authority remains accountable, and automation is possible only under the applicable legal conditions.

2.3 Functions and life events describe different views of the same system

From the citizen's perspective, public administration is often experienced through life events such as starting a business, becoming unemployed, having a child, or dealing with a bereavement. A single life event may involve several organisations and services. Yet beneath those different services, many of the same functions recur: identification, application intake, assessment, decision, notification and payment.

Life events describe the person's situation. Functions describe the recurring public capabilities needed to respond to it. The two perspectives therefore complement rather than replace each other.

Next right questions

Short term – next 18 months	Long term – towards 2030
Can we describe one real service from public purpose to task and identify where an agent could contribute, where human judgement is required and where statutory authority applies?	Should our service catalogues, architecture and performance measures use the same function definitions so that shared work becomes visible across organisations?

3. Which functions should be prioritised?

Making recurring work visible does not yet tell leaders where to act. A function may be technically well suited to agentic AI and still be a poor investment. The volumes may be small, the expected benefit limited, the risks disproportionate or the participating organisations unable to deliver the change.

Conversely, a high-value function may be worth redesigning even if only part of the work can be supported by AI. The starting point should therefore be value, not automation potential.

The First Moves assessment helps identify where recurring activities appear and where agentic AI may be technically and legally feasible. That is an important first filter, but not a sufficient basis for investment. Leaders still need to assess the value, scale, deliverability and public-sector readiness of the same work.

For the next 18 months, the recommended starting point is a practical cross-organisational demonstrator: a bounded test of one recurring function. This kind of test should begin only when leaders can answer four connected questions about the same function.

These are not four separate assessments. They must describe the same function. A strong case emerges only when all four perspectives point to the same opportunity.

This sounds obvious, but public organisations often describe the same work differently depending on whether they are assessing cost, service value, operational feasibility or technical readiness. This makes it hard to judge the opportunity as a whole: one team may know that a task can be automated, while another may hold the cost, service or organisational information needed to decide whether the wider function is worth redesigning.

Table 2. Four questions for selecting a practical cross-organisational demonstrator.

View	Leadership question	Minimum evidence
Value	What would improve, for whom, and what value would it create?	A clear benefit hypothesis covering cost, capacity, speed, quality, service outcomes and rights.
Scale	Is there enough recurring work for the improvement to matter beyond one local process?	Volumes, staff time, waiting time, current cost and evidence of repetition across services or organisations.
Feasibility	Which parts of the workflow could an agent safely and legally support?	Workflow map, human–agent boundary, legal constraints, data availability and critical failure modes
Deliverability	Can the participating organisations actually redesign, govern and sustain the change?	Named leadership, implementation capacity, partners, funding logic, and a credible route to adoption or exit.

Function-based thinking creates a common reference point across these views. For example, an application-intake demonstrator should not begin simply because intake contains repetitive tasks that an AI agent could perform. Leaders should first ask whether the function contains enough avoidable delay, manual work, incomplete submissions or repeated requests for information to justify intervention.

They should then establish how widely the same problem recurs. If several organisations perform substantially similar intake work, the potential value may extend beyond one service or one organisation.

Only then should the question become more technical: Which steps can an agent support, which information can be retrieved from authoritative sources, where does professional judgement begin, and what legal or operational boundaries must not be crossed?

Finally, leaders need to know whether the participating organisations can turn the opportunity into actual change. A technically strong solution creates little public value if workflows remain unchanged, local systems cannot connect to it, responsibilities are unclear or no organisation has the capacity to adopt it. The ability to deliver is as important as the opportunity itself.

3.1 A shared language makes comparison possible

To compare functions across organisations, leaders need sufficiently consistent definitions. Costs, workflows, benefits, risks and technology must refer to the same recurring work, even when different organisations currently describe that work in different ways.

The same problem appears in data. Two systems may technically exchange information while the organisations behind them interpret that information differently. Semantic interoperability is critical because organisations must agree on the operational and legal meaning of the information they intend to reuse.

The goal is not to create a complete national classification before practical testing can begin. For the first demonstrator, the requirement is simpler: describe one function consistently enough that its value, scale, feasibility and deliverability can be assessed together.

For application intake, this would mean establishing the present volume and delay, the main sources of unnecessary work, the expected value of reducing incomplete submissions, the legal and service risks, the organisations that could reuse the capability and the evidence that would justify continuing, changing, scaling or stopping the test.

A high readiness score is a good start, but a strong priority case should also show why the function is worth prioritising, how large the opportunity is, where agentic AI can genuinely contribute, and whether the public sector can turn the opportunity into measurable change.

Next right questions

Short term – next 18 months
Which potential demonstrator has the strongest combination of value, scale, feasibility and deliverability, and what evidence is still missing before we invest?

Long term – towards 2030
How should Finland maintain shared definitions that allow functions to be compared consistently across costs, outcomes, organisational capability and risk?

4. What must change before AI can improve productivity?

AI does not create productivity simply by making an existing task faster. Productivity improves when technology allows public organisations to remove unnecessary work, reduce handovers, reuse information, redesign roles and deliver the same or better public outcome with less total effort. The leadership challenge is therefore to redesign the work, not simply automate it.

The sequence is crucial: first redesign the work, then agree on the data and responsibilities around it, and only then decide how systems and AI should support it. If technology is introduced before these questions are answered, it can become another layer on top of the existing process. The organisation gains a new tool but keeps the old work, controls and information flows, allowing local efficiency to improve while overall productivity barely changes.

Let us look at application intake as an example once again. Before asking which AI solution should be used, leaders should ask:

- Do we need every piece of information we currently request?
- Does another public authority already hold some of it?
- Which checks are genuinely necessary?
- Which steps could be shared across services or organisations?
- Where is professional judgement required?
- Which handovers, requests or controls could disappear altogether?

The largest productivity gain may come not from automating a step, but from removing the need for that step.

4.1 Data must be understandable, not merely available

Agentic work depends on data, but three distinct questions are often treated as one. Is the information available digitally? Can systems exchange it? Do organisations mean the same thing by it? Table 3. elaborates on the significance of these questions.

Table 3. Three essentials for data reuse: availability, exchange and shared meaning

Question	What it means	Why it matters
Is the information available digitally?	The necessary data exists in usable form.	Without digital access, an agent cannot use it reliably.
Can systems exchange it?	Information can move between systems through compatible interfaces and structures.	AI can help translate formats, but access, security and ownership still need to be governed.
Do organisations mean the same thing by it?	Concepts and values have a shared operational and legal meaning.	An agent can process data, but it cannot legitimately decide what *applicant*, *address* or *complete application* should mean across public authorities.

Two systems may technically exchange information while the organisations behind them interpret that information differently. This is why *semantic interoperability* is critical. It means agreeing on the meaning of the concepts being shared, not only on the technical format in which they travel. Shared meanings are not merely a technical detail. They are part of how the function is governed.

For the first demonstrator, Finland does not need a complete national semantic model. It needs something much smaller and more practical: a governed vocabulary for the selected function. The participating organisations should agree on the meaning of the most important concepts, identify the authoritative source for each key data item, name its owner, record where the information came from and define how errors can be corrected.

For application intake, this might mean agreeing what counts as a *complete application*, what information can be retrieved from an authoritative register, when information must be requested from the applicant and which organisation is responsible for correcting an inaccurate record.

These definitions are not documentation added after the system is built. They are part of how the human–agent system is controlled.

4.2 Build around durable work, not temporary systems

Public organisations often face an architectural dilemma. Replacing a core information system may create a durable foundation, but major system renewals can take years. Waiting for every legacy system to be replaced would make experimentation with agentic models unnecessarily slow.

A more practical route for the next 18 months may be to create stable interfaces around agreed function concepts while existing systems remain in place. The distinction is important because functions should outlast individual systems.

If a new capability is designed around the terminology and structure of one legacy application, it may need to be rebuilt when that system changes. By contrast, if it is designed around more durable concepts — such as applicant, application, authoritative information, completeness, routing and review — individual systems, AI models and suppliers can change without redefining the public work itself.

Technology should remain a replaceable component within a governed operating model. Procurement should follow the workflow, data meanings, safeguards and desired public outcome. It should not determine them.

4.3 Human responsibility must be designed into the workflow

The same function can contain very different kinds of work. An agent may be able to retrieve information, check completeness, compare data against explicit rules, prepare an output or route a case. Other steps may require professional judgement or the exercise of public authority.

The boundary should not be drawn around the function as a whole. It should be drawn inside the workflow.

For every material step, leaders should be able to answer:

- What may the agent do?
- When must it escalate?
- Who remains responsible?
- How can a person understand, challenge or correct the result?

For application intake, an agent might identify that information appears to be missing. But the applicant must be able to understand what is being requested, correct inaccurate information, and reach a responsible person when necessary. These safeguards are not constraints added after productivity has been achieved. They are part of the design of a trustworthy and reusable capability.

Where an action affects a person's rights or the exercise of public authority, the applicable statutory requirements for responsibility, automated decision-making, documentation and remedy continue to apply (Finnish Constitution 731/1999; Administrative Procedure Act 434/2003; Act on Information Management 906/2019).

4.4 Shared capability can also scale failure

Reusing the same capability across several organisations creates productivity potential, but it also introduces shared risk. A weakness in one organisation's process affects that organisation; a weakness in a shared component may reproduce the same error, unequal outcome or vulnerability across many organisations at once. Monitoring, traceability and safe stopping mechanisms must therefore be part of the design rather than later additions (Autio et al., 2024).

Before reuse of capability is scaled, participating organisations need to agree on:

- human review and escalation
- explanation and correction
- appeal and remedy
- privacy and data minimisation
- information security and cyber resilience
- monitoring for unequal or unexpected outcomes
- logging and auditability
- a safe way to pause or stop the capability

A shared capability should be scaled only when both the capability and its controls have worked together in practice.

4.5 Five gates before technology becomes a productivity investment

The first demonstrator should pass five connected gates. They are not intended to create a heavy approval process. Their purpose is to prevent the public sector from investing in technology before the underlying productivity logic is clear.

1. **Public outcome and work removed**
   What public outcome should improve?
   Establish the current baseline for cost, effort, delay, quality, errors and rights. Then identify which work will be removed, combined or stopped if the redesign succeeds. A proposal that only makes existing work faster has not yet demonstrated a productivity gain.

2. **Shared meanings and authoritative information**
   What information does the function actually need?
   Agree on the minimum necessary concepts, the authoritative source and owner of each key data item, how provenance is recorded and how incorrect information can be corrected.

3. **Legal and human boundary**
   What may an agent do — and what must remain with a person or public authority?
   Map statutory competence, agent actions, escalation points, human review, explanation, correction and routes for remedy.

4. **Replaceable architecture**
   Can the technology change without forcing the public work to be redesigned again?
   Require interfaces around durable function concepts, modular components, logging, monitoring, data portability and credible supplier exit.

5. **Evidence and stop conditions**
   How will leaders know whether to continue? Measure time and cost alongside service quality, outcomes for different user groups, staff experience and rights impacts. Agree in advance what evidence would trigger a change, pause or stop.

A technology demonstration that cannot answer these five questions may still be useful. But it is a technology trial, not yet a productivity demonstrator.

Next right questions

Short term – next 18 months
Which workflow should we redesign before procurement — and which meanings, human boundaries, safeguards and stop conditions must all participating organisations agree first?

Long term – towards 2030
Which core systems need renewal, where can stable interfaces around shared function concepts provide a safe bridge, and who should maintain those meanings and standards over time?

5. How can the Agentic State learn across organisations?

Redesigning a function creates only the first opportunity. The larger question is whether evidence from the new human–agent system can improve the same recurring work over time and across organisations.

Exceptions, incorrect information requests, human overrides, delays, complaints or unequal outcomes can all reveal weaknesses in the workflow behind them. If several public bodies recognise that they perform the same function, evidence from one implementation can reveal a weak workflow, unclear data definition, unnecessary control or recurring exception that may also exist elsewhere.

The Agentic State becomes a *Learning State* when evidence from human–agent work is used to improve workflows, roles and rules, not only the performance of the AI.

Today, that learning often stops too early. Public organisations typically govern projects vertically: A project team may be able to adjust the technology or local workflow, while decisions about roles, systems, funding, rules or organisational responsibilities sit elsewhere. A successful experiment can end without changing the wider operating model.

Learning must mean more than improving the model or documenting results. When the same problem recurs, leaders should ask what in the surrounding system needs to change. If an agent repeatedly asks for information another authority already holds, the answer may be to change the data flow rather than the prompt.

If professionals repeatedly override the agent at the same point, that may reveal where human judgement should explicitly remain in the operating model. And when one organisation solves a recurring problem, looking at the work as a shared function makes it possible to ask whether the same lesson should change the function elsewhere.

5.1 Four leadership responsibilities

This changes the role of leadership. Leaders are not only sponsors of AI adoption. They become responsible for creating the conditions in which the human–agent system can learn.

Four responsibilities are particularly important.

1. **Set the purpose and boundaries.** Define what outcome the function should produce, which rights and public values must be protected, and what an agent must never do without human involvement.

2. **Design the human–agent system.** Decide deliberately what work agents can perform, where professional judgement adds value and which exceptions must be escalated. People remain responsible for decisions that require human or legal authority.

3. **Turn evidence into decisions.** Make material errors, delays, overrides, complaints and feedback traceable to a decision: What should change, who decides and when will the effect be reviewed?

4. **Create a route from local learning to wider change.** Ensure that repeated evidence can influence not only a task or workflow, but also systems, roles, funding, rules or oversight when necessary.

Without this route from local learning to wider change, the public sector may become very good at running AI pilots while remaining poor at institutional learning.

5.2 Learning needs to travel through three levels

Evidence from human–agent work can lead to change at different levels. Organisational learning occurs when experience changes shared routines, while deeper learning may require questioning the assumptions, roles or rules behind those routines (Crossan et al., 1999; Argyris & Schön, 1978). The practical question for shared functions is how evidence moves from one level to the next.

Research on network learning adds an important point: organisations can also learn collectively across boundaries when they develop shared interpretations and change common practices (Knight, 2002).

Operational learning improves the immediate work. A task, workflow, instruction or control changes because evidence shows that something can be done better.

Organisational learning changes the capabilities around that work. Roles, systems, skills, responsibilities or management practices may need to change.

Institutional learning addresses barriers that no individual organisation can remove. Repeated evidence may point to a need to change shared standards, funding arrangements, mandates, regulation or oversight.

A demonstrator has system-level value only if there is a route for evidence to travel through these levels when necessary. This does not mean that every operational problem should become a national policy question. Most should not. The point is to create a clear escalation path for problems that recur across organisations or cannot be solved within the existing operating model.

5.3 What should be learned from the first demonstrator?

For an application-intake demonstrator, success should be measured through more than processing time or staff hours saved.

The participating organisations should examine:

- incomplete applications and unnecessary requests for additional information;
- incorrect or inconsistent evidence requests;
- human overrides and recurring exceptions;
- delays and handovers;
- complaints, correction requests and unequal outcomes;
- staff and user experience; and
- whether learning from one organisation changes the shared function elsewhere.

For every material signal, the demonstrator should record a simple decision trail: What happened? Why did it happen? What should change? Who can decide? Did the change improve the outcome? That learning loop is as important as the agent itself.

Two opposite risks need to be avoided:

1. **Experimentation without shared learning** leaves useful evidence trapped inside individual projects.

2. **Governance without learning** creates common rules around processes that may themselves need to change.

The goal is neither unrestricted experimentation nor rigid standardisation. It is a public sector that can learn quickly while keeping responsibility, rights and democratic accountability clear. During the next 18 months, this can be tested without creating a new national governance structure. Participating organisations can agree on a common evidence review, identify which signals require a change decision and make the resulting decisions visible across the demonstrator.

Towards 2030, the larger question is institutional: *Who is responsible for ensuring that learning from shared functions changes standards, capabilities, funding, and oversight when the evidence calls for it?*

Next right questions

Short term – next 18 months
Which shared function should become the first real test of system-level human–agent learning? What evidence should inform decisions to continue, change, pause or stop, and who has authority to make those decisions?

Long term – towards 2030
Who should ensure that learning from shared public functions leads to changes in capabilities, standards, funding or oversight when evidence shows that the operating model needs to change?

6. How can shared functions create reusable capability and a GovTech market?

When the same function appears across several public services or organisations, government faces an important strategic choice. Each organisation can continue to design, procure, and maintain its own version of the same capability, or participating organisations can identify which parts of the work are genuinely common and make that core reusable.

The second option does not mean creating one national application or making different public services identical. It means making repeated public demand visible and separating what can be shared from what must remain specific to each service, organisation or legal context. This distinction is central to the function-based approach.

6.1 Different services can share part of the same work

Consider three very different public services: applying for unemployment benefit, requesting a birth certificate, and acknowledging parenthood. Their purposes, legal rules and outcomes are clearly different, and they should remain distinct services. Yet beneath those differences, some of the operational work may be very similar. The question is not whether the services are similar, but whether parts of the work behind them are.

Each service may need to identify the person, receive information or declarations, check whether required information is present, retrieve authorised data from trusted sources, record where information came from, route the case, maintain logs and audit trails, and ensure accessibility, human review and correction. These recurring elements can form *a reusable capability core:* the shared workflow steps, agreed concepts, interfaces and safeguards that can be reused while allowing different services to remain distinct.

The design principle is therefore not to standardise as much as possible, but to standardise what is genuinely common and make necessary variation explicit. If two services share only identity, record-keeping or a single interface, those may be the only elements worth reusing. Function-based thinking should reveal real commonality, not manufacture similarity where it does not exist.

6.2 From shared work to shared capability

A reusable capability is more than software. It may include common workflow elements, agreed meanings, rules, interfaces, tests, safeguards and documentation that allow several organisations to perform the same recurring part of the work safely and consistently.

This distinction is important for procurement because defining reuse as "everyone uses the same application" can easily create a new form of centralisation and supplier dependency. Defining reuse at the level of *capability* creates more room for variation, competition and innovation. Services can share the same core capability while keeping the parts that are unique to their own context.

A shared function does not make different public services identical. Claiming unemployment benefits, requesting a birth certificate and acknowledging parenthood are different services with different legal outcomes. They can nevertheless reuse parts of the same function.

Figure 2. Distinct public services can draw on a reusable capability core while retaining service-specific variation, supported by multiple interoperable applications and shared public digital infrastructure.

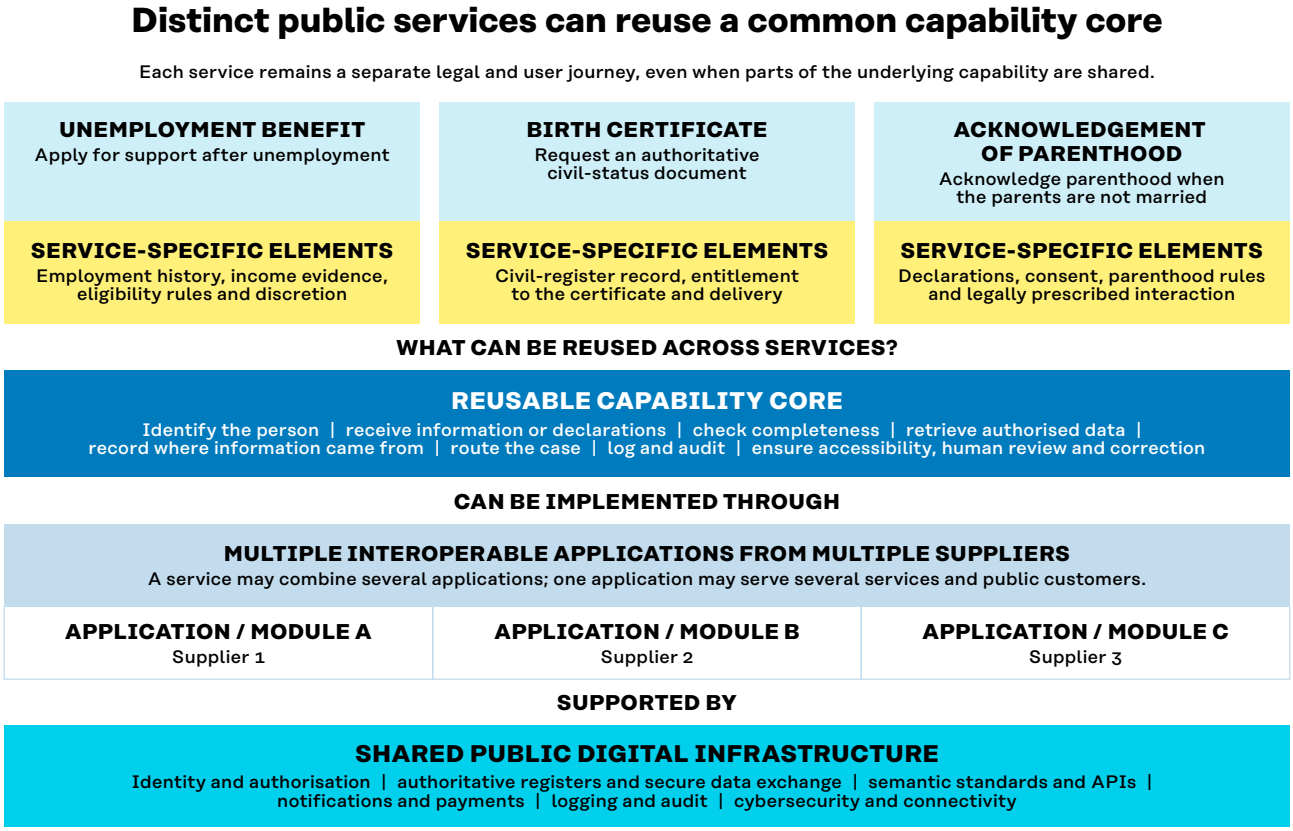

6.3 From shared capability to a GovTech market

The most durable shared asset may be the public definition of the work itself rather than any individual application. It defines what the recurring work must achieve, which concepts and interfaces are common, what safeguards apply and how performance is tested. Individual applications, AI models and suppliers can then evolve around that shared specification.

This is significant because public organisations currently procure many similar capabilities separately. Each describes its needs in its own terminology, procurement process and existing system context. What appears to be a unique requirement from one organisation's perspective may, across the public sector, represent repeated demand for much of the same underlying capability.

Making that demand visible can change the market. If several public bodies agree on the stable common elements of a recurring function, suppliers can develop interoperable applications and modules for several public customers instead of repeatedly building bespoke implementations. Public bodies can share development and assurance costs, while suppliers gain a clearer basis for investing in scalable GovTech products. What appears to be fragmented local demand can become a visible market opportunity once the underlying function is recognised.

Aggregated demand must not, however, become one national system or one dominant supplier. The objective is an open market in which common public definitions, interfaces, testing requirements and safeguards allow multiple suppliers to compete and interoperate. Local and service-specific variation remains

possible, but it should be explicit rather than embedded in incompatible proprietary specifications.

This is why sharing a capability does not necessarily mean sharing an application. Public organisations may perform the same function and reuse parts of the same capability while using different applications and suppliers.

6.4 Shared infrastructure beneath reusable capabilities

Not every common element belongs at function level. Digital identity and authorisation, authoritative registers, secure data exchange, notifications, payments, logging, cybersecurity and connectivity are horizontal foundations that can serve many functions and services.

The model therefore has four layers: distinct public services retain their legal purpose and service journey; a reusable capability core contains only what is genuinely common; multiple interoperable applications and suppliers operate around that core; and shared public digital infrastructure provides common foundations.

The objective is not maximum standardisation. It is maximum useful reuse without losing legitimate variation.

6.5 Existing examples show parts of the logic

Finland's *Incomes Register* demonstrates how a shared data capability can remove repeated work across organisations: income information is reported once and reused for taxation, benefits, insurance, statistics and supervision. In 2025, around 370 organisations used the register in decision-making and income information was distributed approximately 1.7 billion times (Finnish Tax Administration, 2026).

GOV.UK Notify demonstrates a different form of reuse. A bounded government capability for email, text messages and letters serves a large number of services and public organisations without requiring their wider workflows to become identical (GOV.UK Notify, 2026).

Neither of the above is a complete example of function-based government. Together they show the underlying principle: when the common element is carefully bounded, reuse can become the default without requiring services, organisations or suppliers to become identical.

6.6 What would this mean for application intake?

For application intake, the shared capability might cover identifying the applicant, receiving information, checking completeness, retrieving authorised data, recording provenance, routing, logging, accessibility and human review. Service-specific evidence, eligibility rules, judgement and legal interaction would remain separate.

Before procurement begins, participating organisations should therefore be able to answer: What is genuinely common? What must vary? What remains publicly governed? How will components interoperate and be tested? Who maintains the shared assets, and how can an organisation or supplier exit?

When these questions can be answered consistently, a repeated administrative problem becomes both a reusable public capability and a visible market opportunity.

Next right questions

Short term – next 18 months
Which parts of one recurring function are genuinely common across at least three services or public bodies, and what open specification would allow them to be reused safely across multiple applications and suppliers?

Long term – towards 2030
Who should steward the reusable core and shared infrastructure, maintain semantic and technical standards, and ensure that several suppliers can continue to serve several public customers without lock-in?

7. How should shared functions be governed across organisational boundaries?

The function-based approach creates an unavoidable governance question. A recurring function may cross organisational boundaries or appear in similar form across several public bodies, while budgets, statutory authority and accountability remain tied to individual organisations.

This creates a genuine tension. Leaders may recognise the value of redesigning a function together but still need to justify why their organisation should fund the work, accept new risks, or change its processes when some of the benefits may arise elsewhere. Function-based approach will not become sustainable if these questions are treated as implementation details.

At the same time, the answer cannot be vague collective ownership. Public administration requires clear accountability. Someone must remain answerable for decisions, data, services and the exercise of public authority. A function-based approach must therefore add coordination across organisations without weakening responsibility within them.

The challenge is not to assign a new owner for the function. It is to coordinate the shared function sufficiently that participating organisations can improve it together while keeping statutory responsibilities clear.

7.1 Start with a coordination mandate, not new institutional ownership

For the next 18 months, Finland does not need to resolve the permanent ownership of cross-organisational functions before anything can be tested. A more practical starting point is coordination under existing mandates.

Participating organisations can appoint a named convenor, or *function steward*, with a bounded coordination mandate. The steward maintains the shared view of the function, coordinates evidence, prepares joint recommendations and ensures that agreed questions reach the appropriate decision-makers. Statutory authority, organisational budgets and control of individual services remain with the relevant organisations.

Table 4. Coordinating shared functions while keeping organisational accountability clear.

The shared function *Where coordination is needed*	**The participating organisation** *Where accountability remains*
May cross organisational boundaries or recur in several public bodies.	Holds a specific statutory mandate, budget and organisational responsibility.
Creates opportunities, dependencies and risks across the end-to-end work.	Remains answerable for its own decisions, staff, data and services.
Benefits from shared measures and coordinated improvement.	Continues to operate within its own legal and financial accountability.

The steward helps participating organisations see and improve the whole without becoming legally responsible for every part of it.

Such a model makes it possible to test whether a function-based approach produces better decisions before Finland creates permanent structures around it. It also gives leaders evidence about which governance problems can be solved through cooperation, and which genuinely require changes in mandates, funding or legislation.

7.2 Four practical problems must still be solved

Even a bounded stewardship model needs to address four recurring objections.

1. **No one controls the whole.** A shared function may depend on several organisations, systems and legal responsibilities. The convenor therefore needs a clear written mandate to coordinate the work, while the limits of that mandate must be equally explicit.

2. **Costs and benefits may fall in different places.** One organisation may finance redesign or integration while another receives much of the benefit. The demonstrator should therefore record who contributes resources, where benefits are expected to arise and how these effects will be assessed.

3. **Organisations may have legitimate differences.** The fact that a function is recognisable across several organisations does not mean every implementation should be identical. Shared elements should be modular, while justified variation remains possible where legal, operational or service reasons require it.

4. **Going first creates risk.** Early participants invest management attention, staff time, legal work and implementation capacity before the wider value is known. They therefore need clear cost-sharing arrangements, an agreed exit route and recognition of the capacity they contribute to building a potentially reusable capability.

These are not arguments against the function-based approach. They are governance conditions that need to be made explicit so that cooperation remains rational for every participant.

7.3 What can be done within the next 18 months?

A first demonstrator can operate largely within existing mandates if the participating organisations agree on a small number of governance arrangements before development begins. They should:

1. Select one shared function and agree why they are working on it together.

2. Sign a short participation agreement that records the purpose, responsibilities, contributions and limits of the collaboration.

3. Appoint a named function steward with a mandate to convene the work, maintain shared evidence and prepare recommendations.

4. Establish a common baseline and at least one end-to-end measure that allows the participating organisations to assess the function as a whole.

5. Map legal responsibility, data roles, human review, escalation and routes for remedy before technology is introduced.

6. Agree how shared costs, local variation, supplier access and exit will be handled.

7. Review evidence at predetermined decision points to continue, change, pause, stop or scale the work.

These arrangements do not override the organisational structure of public administration. Each organisation continues to make the decisions for which it has legal and financial responsibility. What changes is that some of the evidence, design work and recommendations are created together rather than repeatedly in isolation.

The first governance model is therefore deliberately bounded. Its purpose is to test whether organisations can coordinate one recurring function well enough to learn what shared governance actually requires before any permanent structures are created.

7.4 Governance must follow the level of the decision

The five-level model introduced in Section 2.1 helps clarify which decisions belong where.

Public purpose and major portfolio priorities require political and senior administrative leadership. Decisions about statutory mandates, public funding and accountability remain with the institutions legally responsible for them. Workflow design belongs close to the organisations and professionals doing the work.

Function level has a narrower but important role. It allows participating organisations to compare performance, identify shared dependencies, maintain common definitions, coordinate reusable capability and bring cross-organisational problems to the level where they can be resolved.

This prevents two opposite mistakes. One is trying to centralise decisions that should remain close to organisations and professional practice. The other is leaving genuinely shared problems to voluntary coordination even when no individual organisation has the incentive or authority to solve them.

7.5 Towards 2030: Some questions will require structural decisions

An 18-month demonstrator can test function-based coordination, but it cannot settle the permanent governance of shared public capabilities. If selected functions prove to create significant cross-organisational value, Finland will eventually need structural decisions on four issues: mandate, funding, leadership, and standards and oversight.

1. **Mandate**: Which functions are important enough to justify formal coordination across organisations, and who provides that mandate?

2. **Funding**: How should shared development, maintenance and adoption be financed when costs and benefits are distributed across several public bodies?

3. **Stewardship**: What permanent stewardship is needed at the shared-function level, and how should it connect with decision-makers inside participating organisations?

4. **Standards and oversight**: Who maintains the common definitions, interfaces, professional capability, assurance and safeguards on which reuse depends?

The UK Government Functional Model provides one useful precedent. It shows that horizontal functional leadership and common standards can coexist with departments and local function leaders (UK Government, 2021–2026). It should not be treated as a ready-made model for Finnish citizen-facing services, particularly where decisions affect individual rights, but it demonstrates that horizontal coordination and vertical accountability can coexist.

The longer-term objective is therefore not to create a new layer of administration around every function. It is to ensure that where public work is genuinely shared, the ability to improve it, is not constrained by organisational boundaries that the work itself repeatedly crosses.

Next right questions	
Short term – next 18 months Who should convene the first shared-function demonstrator, what should that role coordinate, and where do participating organisations remain separately accountable?	**Long term – towards 2030** Which shared functions require a permanent mandate, funding model and stewardship arrangement, and how can these strengthen coordination without weakening clear public accountability?

8. What should leaders do next?

The first step towards a more agentic public sector should not be a national platform or a large technology programme. It should be a bounded test of the central proposition of this working paper: whether looking at recurring public work through functions helps organisations redesign work, reuse capability and learn across boundaries.

Application intake, including initial processing, is a credible starting point because it recurs across the public sector, contains measurable repetition and delay, and can be separated from the final legal decision. The first demonstrator should test the proposition rather than assume it is true.

8.1 Decide now

A senior leader in one public body should invite at least two other public bodies that perform substantially similar application-intake work into a short scoping phase.

Together, they should describe one real service from public purpose to task and establish the present baseline. They should identify where professional judgement and legal authority begin, appoint a function steward and agree what evidence would justify continuing, changing or stopping the work.

The first decision is therefore not which AI solution to procure. It is whether several organisations can describe the same recurring work clearly enough to decide what should be redesigned together.

8.2 Prepare next

The participating organisations should then redesign the workflow before choosing technology. They should identify work that can be removed or simplified, information that can be reused, the smallest set of shared concepts, the human–agent boundary, and the parts of the capability that are genuinely reusable.

Technology should then be tested against the five gates described in Section 4.4: public outcome and work removed; shared meanings and authoritative information; legal and human boundary; replaceable architecture; and evidence and stop conditions.

The demonstrator should also create a visible learning loop. Material errors, delays, overrides, complaints and unexpected outcomes should lead to explicit decisions about whether the workflow, data, role, rule or control needs to change — and whether the same lesson applies elsewhere.

8.3 Do not decide yet

One demonstrator is not enough to justify a national platform, a transfer of statutory authority, the standardisation of entire public services or system-wide productivity claims. Nor does a reusable capability imply that every organisation should use the same application or supplier.

Those choices require wider evidence about public value, productivity, rights, service quality, organisational capability, cost allocation and systemic risk. The first demonstrator should therefore be designed to reduce uncertainty, not justify a predetermined national solution.

8.4 Decide towards 2030

If evidence supports broader adoption, selected shared functions may require durable stewardship, common standards, cross-organisational funding and continuing oversight. Shared capabilities may also require

arrangements that keep interfaces open, preserve public control over core concepts and allow several suppliers to serve several public customers.

Function-based coordination should strengthen accountability, rather than blur it. Legal decisions, remedy and responsibility for the exercise of public authority remain with the responsible organisation unless legislation explicitly provides otherwise.

8.5 The practical starting point

The immediate action is deliberately simple: Find one important function that several public bodies already perform. Make that shared work visible. Establish what it currently costs in effort, delay, quality and risk. Then decide what those organisations should redesign, learn or build together.

The first demonstrator should ultimately test two propositions: whether a reusable capability can improve the same function across several public bodies, and whether evidence from human–agent work can lead to changes in workflows, roles, rules or shared capability across organisational boundaries.

If both propositions hold, the significance goes beyond one AI use case. Finland will have begun to demonstrate a new way of organising public-sector renewal: one that starts from recurring public work, uses technology as an enabler and turns shared learning into measurable public value.

Essential terms

Term	Meaning in this working paper
AI agent	Software that can perceive inputs, reason and plan bounded actions, and act through digital tools or interfaces to achieve a defined objective with a degree of autonomy. It can perform several steps towards a goal and escalate exceptions or boundaries it must not cross. It is not a public official or legal decision-maker.
Agentic State	A public administration in which people and AI agents can complete multi-step work under clear human authority, shared rules and public-purpose accountability.
Learning State	A public system capable of turning evidence from human–agent work into changes in workflows, roles and rules and, when needed, institutional arrangements across organisational boundaries.
Function	A recurring public-sector activity that produces a discrete output and consists of at least one end-to-end workflow of interconnected tasks.
Demonstrator	A bounded cross-organisational test of one recurring function, used to assess reusable capability, public value and learning across organisational boundaries.
Public service	The citizen- or business-facing service delivered within a particular country's or region's legal and institutional context. One service can use several functions; one recurring function can support several services.
GovTech market	A competitive market in which a reusable capability core, open interfaces and shared infrastructure allow multiple applications from multiple suppliers to serve several public customers.
Workflow	The sequence through which a function produces its result; the level where work is redesigned and the human–agent boundary is set.
Semantic interoperability	Agreement on what data and concepts mean across systems and organisations.
Shared public digital infrastructure	Common foundations—such as identity, authoritative registers, secure data exchange, semantic standards, notifications, logging and cybersecurity—used by many services, applications and suppliers.
Shared capability	A reusable core of agreed workflows, meanings, interfaces and safeguards that supports the common parts of a recurring function while allowing explicit service-specific variation.
Function steward	A named convenor with a bounded coordination mandate for maintaining the shared view of a function, coordinating evidence and preparing joint recommendations. Statutory authority remains with the responsible organisations.
Systemic productivity	Improvement in public outcomes or capacity relative to total effort, achieved by removing avoidable work, delay and duplication across organisations rather than only accelerating an individual step.

References

Act on Information Management in Public Administration 906/2019 (Laki julkisen hallinnon tiedonhallinnasta 906/2019), Chapter 6 a. https://www.finlex.fi/fi/lainsaadanto/2019/906

Administrative Procedure Act 434/2003 (Hallintolaki 434/2003), Chapter 8 b and sections 53 e–53 g. https://www.finlex.fi/fi/lainsaadanto/2003/434

Agentic State foundation & Sitra (2026). *First Moves: A Finland-Specific Assessment of Agentic AI Potential in Government*. https://agenticstate.org/projects/finland. See also the First Moves Dashboard: https://firstmovesinteractivedashboard.sitra.fi/

Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.

Autio, C. et al. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

Crossan, M. M., Lane, H. W., & White, R. E. (1999). *An organizational learning framework: From intuition to institution*. Academy of Management Review, 24(3), 522–537. https://doi.org/10.5465/amr.1999.2202135

Finnish Constitution 731/1999, especially sections 21, 22 and 124. https://www.finlex.fi/fi/lainsaadanto/1999/731

Finnish Tax Administration (2026). *Incomes Register – About us and 2025 statistics*. https://www.vero.fi/en/incomes-register/about-us/

Ilves, L. et al. (2025). *The Agentic State: Rethinking Government for the Era of Agentic AI*. Tallinn Digital Summit. https://agenticstate.org/paper.html

Knight, L. (2002). *Network learning: Exploring learning by interorganizational networks*. Human Relations, 55(4), 427–454. https://doi.org/10.1177/0018726702554003

UK Government (2021–2026). *Guide to Functional Standards* and *Implementing Functional Standards: A Guide for Senior Leaders in Organisations*. https://www.gov.uk/government/publications/guide-to-functional-standards/guide-to-functional-standards

UK Government (2026). *Performance*. https://www.notifications.service.gov.uk/features/performance

World Economic Forum, Capgemini & Global Government Technology Centre Berlin (2026). *Making Agentic AI Work for Government: A Readiness Framework*. https://www.weforum.org/publications/making-agentic-ai-work-for-government-a-readiness-framework/

SITRA

SITRA WORKING PAPER September 2026
Sitra working papers provide multidisciplinary information about developments affecting societal change. Working papers are part of Sitra's future-oriented work conducted by means of forecasting, research, projects, experiments and education.

ISBN 978-952-347-483-3 (PDF) www.sitra.fi

SITRA.FI
Itämerenkatu 11–13
PO Box 160
FI-00181 Helsinki
Finland
Tel: +358 294 618 991
Sitra