

A Government-wide Enterprise Architecture

How To Avoid Indigestion When Attempting To Eat An Elephant

Preamble

This paper outlines the route map to a Government-wide Enterprise Architecture. It covers the aims, likely issues, potential components and next steps in achieving delivery.

The words “Enterprise Architecture” particularly delight technologists as there is a tidy split between business, the “enterprise”, and technology “architecture” – both words, however, are relatively dry and it is difficult to conceptualise what the output of such a thing might be. It is perhaps best to think of the Enterprise Architecture as set of overlapping blueprints covering the flow of tasks through business units underpinned by technology (and here, all aspects of technology are included, such as, networks, systems, data and so on) – but to think that this is a technology initiative, led by the IT department, will result in failure.. The idea of an Enterprise Architecture has been around since 1987 when Zachmann (www.zifa.com) first proposed the idea and conceived a series of frameworks that would help model them. Today, the concept is mature with a large number of consultancies offering assistance in scoping the work, mapping the processes and recommending forward strategies.

For simplicity, enterprise architecture is here shortened to EntArch, pronounced EntArk; also, frequent mention is made of government departments or just departments – at all times, this shorthand should be taken to include all government entities, whether they are central departments, agencies, local authorities, the NHS, NDPBs or whatever.

For government, an EntArch is very much a “you can’t get there from here” problem. Historically, efforts have been few and far between in implementing such a framework and very few departments are exploring their business processes in anything like the detail or with anything like the kind of cold and calculating eye that would institute the change necessary to realise the benefits of an EntArch. Almost every department has a “one of everything” architecture with a myriad of overlapping processes & systems, bundles of duplicate data and inconsistent presentation to the customer. In many ways, that is an accident of how technology has matured rather than particularly the fault of a department, individual or supplier.

Recently though, plans have shown signs of change and some of the larger departments are at least embarking on the technology side, even if the business re-architecting is not yet happening at the same pace. To deliver a full EntArch though, business and technology must be marching in lock-step – else redundant processes are automated, the business is put through the pain of unnecessary transition (or poorly thought out IT projects that hinder the customer service process) and strategic benefits will not be achieved.

What's driving the focus on such architectures across both public and private sector organisations is a recognition that IT isn't working – it's failing to deliver the kind of benefits that were expected and it isn't supporting the business. The Meta Group note that only 25% of projects deliver on their narrowly-defined "project only" goals and only 12% deliver any strategic business advantage. At the same time, business units are introducing new products with equivalent, but different, processes to existing products leading to duplication in both business and technology.

Also, for too long technology implementations have been viewed on a per-project basis without recognising the case for a consistent and disciplined use of technology throughout an entire organisation. *Project successes are today viewed as heroic*, but must become systemic for the organisation to have any real chance of adapting to change.

The main candidate for how to achieve this is to promote reuse on an organisation-wide basis so that successful projects feed other projects ... all boats float. Newly emerged technologies can be leveraged to speed the process, without needing to rip out old technologies in a wholesale fashion – so change can be incremental and self-funding rather than dramatic and capital intensive.

The output of a successful EntArch strategy definition exercise will, at a business level, be streams of horizontally layered processes that are relatively consistent, maybe even common, first within and then across departments. At a technical level, a programme will emerge to move away from the existing, closely-coupled, inflexible systems to a more componentised architecture that can be rapidly deployed, updated and reconfigured to support changes in business goals. Alongside business and technology will be new definitions and applications of data standards and organisational models. Underpinning all of this will be sets of indicators, vocabularies and measures that ensure it all stays on track.

We must take on the lessons learnt as we have tried to deliver central infrastructure for e-government and find a new way to deliver an EntArch, one that is more likely to succeed. This cannot, however, be done without substantial changes in organisation, funding and collaborative working practices. It will, without doubt, be an arduous and painful process but one that ultimately can succeed if it is given the right executive focus, from a business, technology and financial standpoint.

The bottom line is that departments are going to have to work together both in business process and technology delivery to make this happen. That's not something that we've seen a lot of to date. Gaining buy-in to this enormous task will be challenging and fundamental to any chance of success.

The Aims and Consequences of e-Government

When the programme of e-government was first imagined, few would have disagreed with the following overall aims:

- A resolute focus on the citizen and business owner, eliminating the tendency for interaction through individual silos
- Increased provision of services through a variety of intermediaries helping to eliminate the digital divide
- A substantial reduction in the burden of dealing with the state for citizens and businesses
- The e-channel eventually being established as the preferred channel for many people allowing waste and bureaucracy to be eliminated and significant cost savings to be made

However, it is difficult to see the progress that we have made towards those aims if government is looked at in the whole. If they were the aims, what progress have we made?

- Nearly 3,000 individual government websites with arbitrary and confusing domain names, an enormous duplication of information, a massive variety of navigation plus accessibility standards and a resolute refusal by departments to join up and present services consistently, coherently and in the best interests of citizens.
- A few transactional services each of which is presented in its own environment with limited ability to exchange or reuse data, few chances to use intermediaries to prepare and manage the submission of transactions and a dramatic and internationally visible lack of usage of the few services that exist.
- A Central Infrastructure that meets the needs of some departments but certainly not the majority that has varying degrees of usage by even the departments that ardently support the need for it but is seen by many as inflexible and requiring them to cede control of the processes that they own.
- Major initiatives, such as that for Tax Credits or, for that matter, the PRO 1901 Census Online project, being implemented with widespread negative PR impact following apparent technology issues without any process for other parts of government to identify and learn from the issues

Nowhere to be seen, despite laudable efforts by some, is the reduction in burden, the elimination of bureaucracy or the delivery of outright cost saves – whether they be realised through reuse savings at a pure technology level or organisational savings through reduced use of existing channels. In a culture that promotes the drawing of curtains to hide from others what is really going on, it is difficult to see find ways to

improve. There are lessons to be learnt across all of the business and technology projects handled to date.

Certainly things will have to be different if we are to embark on a further joined up initiative, one that is likely more complicated and certainly more risky in the initial stages than anything that has been done to date in UK government.

What is an Enterprise Architecture?

Definitions vary, as is so often the case with “buzzwords” ... here are a few:

From the US Government’s Clinger-Cohen Act, “Information technology architecture means an integrated framework for evolving or maintaining existing IT and acquiring new IT to achieve the agency’s strategic goals and information resources management goals.” Although useful, this definition doesn’t capture the need for an Enterprise Architecture to encompass both business and technology change.

The Open Group (www.opengroup.com) goes a little further and breaks out the components. Business architecture: the business strategy, governance, organisation and key processes; Data architecture: the structure of logical and physical data assets and data management resources; Application architecture: a blueprint for individual systems to be deployed, their interactions and relationships with/to the business processes; and, Information Technology Architecture: the software infrastructure to support the deployment of core, mission-critical applications – this part is sometimes referred to as middleware.

Ruth Malan, a consultant, defines it as an overall view of all systems in an Enterprise, their properties and interrelationships that plans for and guides the evolution of these systems so that they support and enable the evolution of the business in its pursuit of strategic advantage.

The Metagroup define it as a “top down, business strategy driven process that coordinates parallel streams of internally consistent work involving development of the enterprise, information and technology architectures as well as the delivery of solutions.”

So, the essence of a viable EntArch is that it is an overall view of an organisation, encompassing business and technology processes and systems, both as they are and how they will evolve over time. Principally, *it’s not about technology*, it’s about the intermingling of business goals with technology facilitation. Likewise therefore, it is not especially about delivering government services purely through online channels but about how to achieve an integrated, cross-channel approach that allows citizens and businesses to access services in a consistent manner, no matter the channel that they choose.

The EntArch will change dynamically as short-term objectives are realised and greater benefits are unlocked – so the EntArch is not (yet another) project but a programme that

must be robustly managed, sponsored (*governed*) at a senior level and communicated throughout the organisation.

Government, however, is far from being just another organisation. It is, if anything, an über-organisation – an organisation of organisations with both willing partners and warring federations inside each unit and across several units. What a challenge!

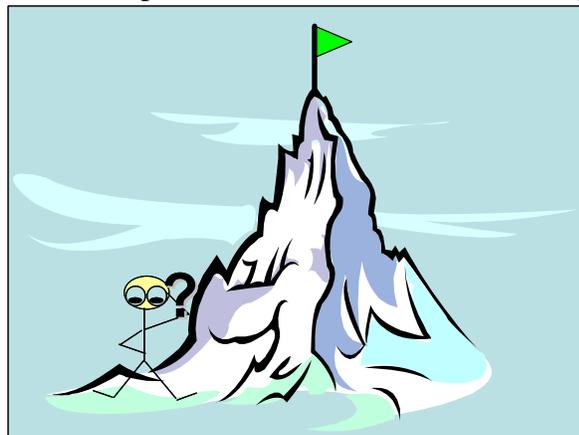
The Aims of an Enterprise Architecture

- To speed the deployment of customer-focused services, facilitating the joining up of otherwise discrete departmentally-managed services
- To provide for widespread service offerings through large varieties of intermediaries, each able to offer different and competing value-added enhancements to any service
- To ensure that lessons learnt in any one organisation in the implementation of any part of this or any other EntArch are passed to all other organisations through providing a collaborative learning environment
- To reduce the cost of delivery and maintenance of services through reuse of systems, components and/or process changes and through rationalisation of the overall number of such systems, components and processes
- To buy time (and provide funding) for the eventual and full rationalisation of departmental back end systems through constructing a flexible and capable integration layer covering the über-organisation
- To facilitate a dramatic and non-linear change in the perception of government service by citizens and businesses.
- Ultimately, to deliver services that citizens *want* to use, that are consistent in their operation and are easy to use, no matter the chosen channel.

You Can't Get "There" From "Here"

First off, understanding where "here" is can be quite a challenge. Reviewing the history of how we got "here" may be useful when considering what needs to happen next.

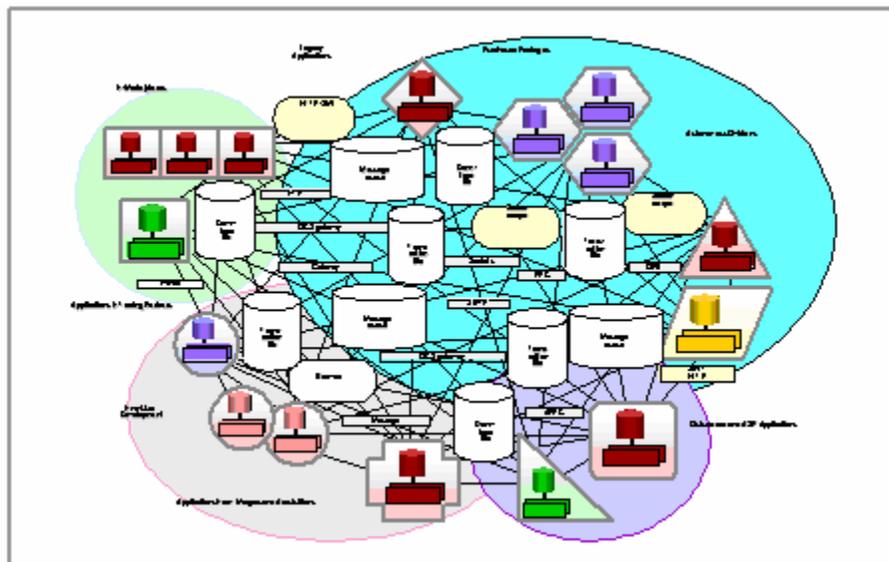
Taking the automation of a general ledger as an example (it's not hard to extend thinking from there to a bank account and then to the "citizen's account" with government), initially all of the entries were made by hand using a double-entry technique – reconciliation (and so data accuracy) was inherent in the process. Periodic audits were held to make sure that the numbers were good and did indeed correspond to actual



transactions. The data was pretty much the be all and end all of the process. *The data was the asset.*

As the general ledger was automated, systems were built on top of the ledger to post new entries. Other systems were built to create reports, manage new projects and enhancements were made to existing systems to provide more functionality. Then client systems were added to track orders, follow up questions, identify potentially fraudulent transactions and so on. Every system had its own copy of the original data, but not necessarily in the same format. Pretty soon, the data was no longer the real asset – so many changes were being made to it by so many systems and there were so many overlapping, duplicate or inaccurate copies that it became almost impossible to introduce meaningful changes in a reasonable period. The new asset was the knowledge of key, long-serving staff in how the systems interacted and treated the data.

The system architecture for such an environment, i.e. “here”, looked like this:

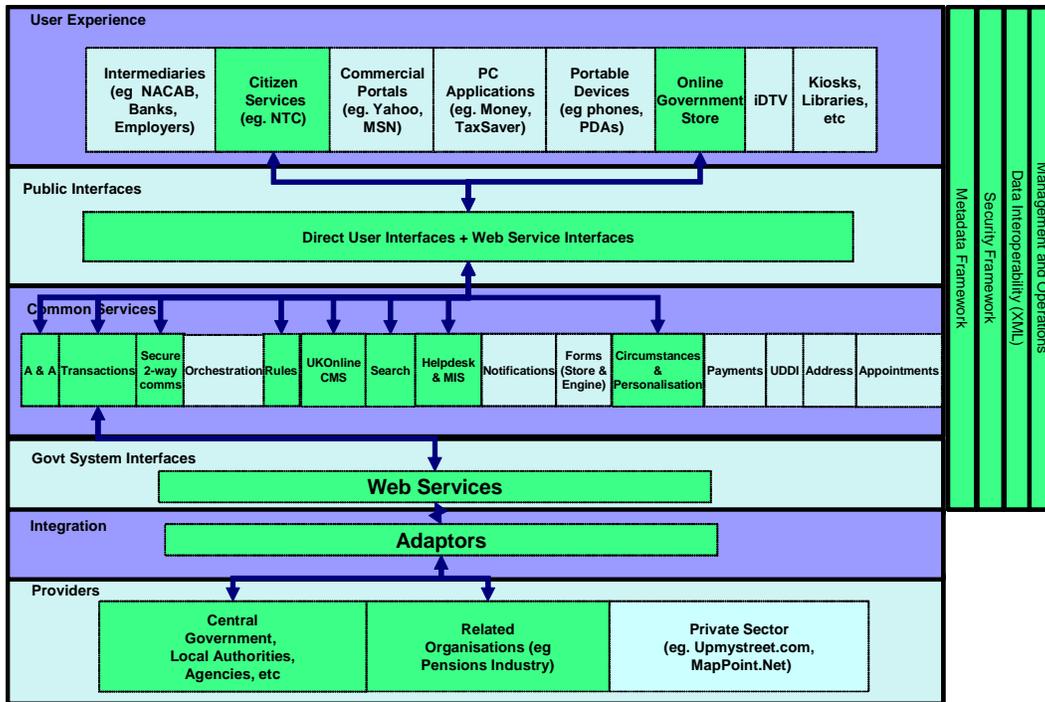


Source: Grazyna Witkowska

This is, of course, Gordius’ knot from the time of Ancient Greek – an intricate and apparently unsolvable puzzle - short of Alexander the Great coming along and cleaving the knot with his sword at least.

Such an architectural picture is not uncommon. The largest government departments have several 10s or even 100s of systems dealing with what, to the outside world, might appear to be one process (VAT, say, or Pensions). The challenge for the Enterprise Architecture work is how to unravel these systems and align them better with business processes so that the aims highlighted above can be achieved.

A model for a Government Enterprise Architecture, or “there”, might look like this:

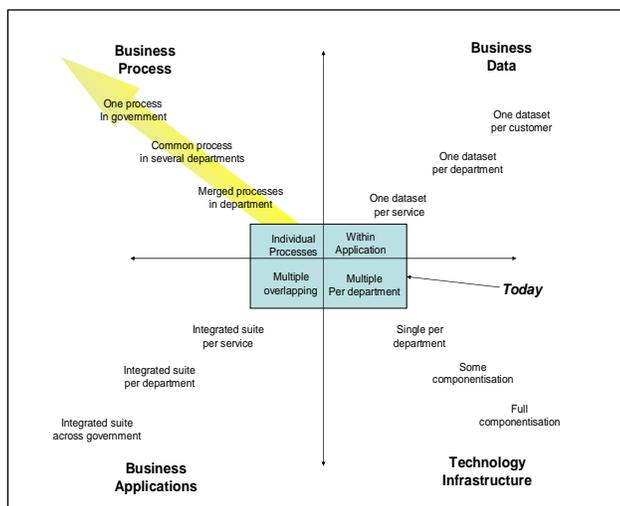


Source: Jerry Fishenden

The principal now is that data can be exposed to many viewers – internal staff, third parties, intermediaries and the citizen/business themselves. The number of data sources has been dramatically reduced, perhaps not to one but to a few at least. This has been achieved principally through abstracting the original back end systems using clever technology known as web services and through creating a set of consistent and reusable components.

The journey to such an enterprise architecture is lengthy – even achieving such a vision in a single department is a huge challenge. It may be appropriate to think of progress being made along 4 axes, not necessarily with equivalent speed. The axes are business process, business application, business data and technology infrastructure - note that the focus is on business involvement and leadership, especially around such important areas as data.

The model might look like the figure at right. Progress is made by moving out along any of the axes, with the time to make progress and the potential for cost saves increasing the further out you move. Although, progress need not be equivalent against each axis at the same time, there will be points when the next level of change can only be achieved when



enough, dramatic progress has been made across each axis.

The Limits of *this* Enterprise Architecture

It would be relatively easy to take on the aims of an EntArch and produce a far-sighted, dream-world view of what government could and should look like, from top to bottom and with all strategic imperatives encompassed. Such a vision would not have much practical use (although there may be merit in holding at least a loose vision of this in mind as successive more realistic iterations are produced).

In this paper, the aims have been toned down to increase the probability of success through reducing the amount of interference in departmental back end processes and focusing instead on the citizen facing view of things as seen from an Internet browser.

In the future, the EntArch must be expanded to cater for such things as common administrative systems (such as HR, Payroll, Finance etc), customer relationship management systems (surely one of the most complex, expensive and risky pieces of infrastructure that any department will undertake) and the upcoming de-layering of back end systems to separate out common components across several parts of our über-organisation (such as receiving and reconciling money, calculating entitlement, the citizen account/government general ledger and so on).

Designing and implementing such an enterprise architecture does, however, closely link with corresponding work on the other functions mentioned. Similar principles will be involved and much can be learned as the initiatives progress. After all, internal staff accessing a system deserve just as much as external customers to access those services consistently and with reliable data presented to them.

The one proviso to structuring this work along the notion of online services is that, in doing so, many problems related to customer data will be encountered and resolved. Systems that then access this data can then be made available both to citizens and also to internal staff – the ability for both the customer (or intermediary) and the civil servant to see the same data will greatly improve efficiency and position the implementation of such systems as customer relationship management tools for much greater success.

How Do We Start?

Embarking on even the definition of an Enterprise Architecture is a significant undertaking that will consume a large amount of resource from across all of government.

Some early decisions need to be taken:

- What is the scope of the initial stage? Trying to define a set of detailed enough strategies across departments that will drive simultaneous change across all of them is not feasible. The initial scope should be limited to some highly visible, high payoff areas – particular areas of benefit delivery (including perhaps both benefits payments

and tax credits – in the past this has been put forward as the “www.giveandtake.gov.uk project”), retirement planning or burden on business (looking at opportunities to rationalise PAYE, VAT, Corporation Tax and so on – some of this may already be going on with the HMT review of Inland Revenue and HM Customs). One of the many consultancies that have experience of EntArch’s could be used to facilitate this process and set up the necessary structures for delivery.

- Who are the players? This is not a job for people who have come to the natural end of their existing job, the so-called “available” people. Departments will have to give up their best and brightest for several months to work on defining and delivering the strategy. The teams should be drawn from business, finance, operations, technology and from inside and outside government (including stakeholders). When working the online solution for Corporation Tax, the Inland Revenue used such an approach (albeit a short, sharp version) to draw out the requirements for the new system.

In addition, the various levels of governance will need to be considered here. Who is coming in from outside? What authority will they have? How will they work with existing teams? How will the programme office be structured? Who will hold the budget? How will decisions be made quickly? Who has the right of veto? *If there is a single characteristic that will define this work it must be charismatic and capable governance.* This programme will test the theory of joined up government to its limits – at a time when there are few other examples of it being seen in the wild on which to base the work.

- How will buy-in be obtained? Allowing this initiative to be voluntary, or adding it to the list of things that departments are already working on will consign it (very quickly) to the rubbish heap of failed programmes. Few things work as well as mandate in the government space. Careful thought will need to be given early on to what is mandatory and what is not and also by when. Setting targets too early in the process will alienate those who need to reach them.

First up, it needs to be mandatory to get agreement on what is going to be done and by when – mandating that it should all be done in a certain way can follow some time later. This should not be allowed to be a lengthy process (no monthly committee meetings here), but one that is fired up and closed down rapidly through a facilitated and focused review programme. Incentives to co-operate later will still be needed – perhaps funding for new projects or new components, perhaps a repayment based on the level of reuse of any one of the department’s own components. Certainly money talks and its use as an incentive will be essential. Other incentives can also be considered though such as giving early participants leading roles in redesigning processes that affect them or giving them more scope in some of the decisions around components and/or vendors.

- What’s the measure of success? Government has had a long history of modernisation, reform and change projects. For this one to succeed, we’ll need to know both the start point and the end point. The earliest session for the senior

sponsors should concentrate on these two key status points and then work a series of milestone with which they would be comfortable.

These milestones can then be refined project by project, once the scope for the whole programme is decided. The teams will need to be clear what their aim is – they will want to know whether it is to streamline a process (making it quicker for the customer), reduce the staff count (saving costs for government), make it more flexible to change (allowing ministers to introduce new initiatives faster) or some combination of these. It's possible that each process has different targets, represented on some kind of sliding scale. Irrespective, they must be clearly spelled out early on so that the teams know what they're working on.

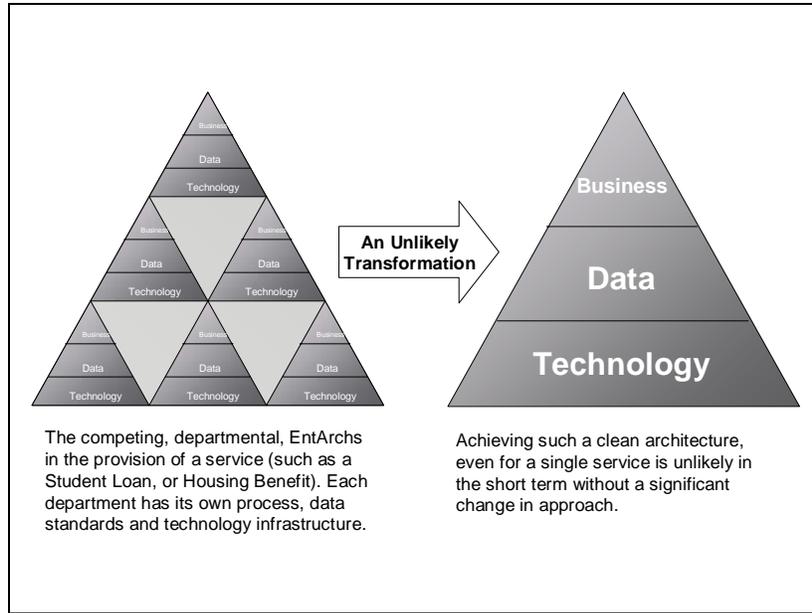
To prevent the milestones being usurped, the target changed or debate about merits, broad sign-up from the community of permanent secretaries and ministerial teams will be needed.

- How will this be funded? Most projects of this nature in the corporate sector are centrally funded through the define and design stages but then rapidly switch to an “as is” basis with ongoing projects saved from eliminating redundant projects in the existing business and technology streams. That is certainly a viable approach here, but will need careful management and process set up to ensure that it is adhered to and that progress does not suffer because of squabbles between departments.

Secondly, within departments, a disciplined approach to funding projects will have to be set up to ensure that they comply with the existing state of the EntArch programme and, where the EntArch standards are lacking, are being used to push forward those missing areas rather than undermine them. The funding process (and its associated governance process) must be flexible enough to ensure that clever and innovative ideas that emerge after the initial budgeting round can still be funded where they fit in with the programme criteria. Too many projects have remained rigid throughout their implementation process resulting in flawed delivery because of seeing every change as a risk rather than viewing the opportunities that may result – this programme, as complex as it is, will need to be fleet of foot to avoid that sin.

The Components of *this* Enterprise Architecture

A single EntArch for the über-organisation is not achievable in a single step – even delivering such a step change for a single, clearly joined-up service is going to take more firepower than has yet been deployed on any single project.



So, rather than get caught up in the politics of control with a vision of getting to “one” EntArch, there is (perhaps) a more pragmatic way to get many of the benefits, avoid some of the pain but still realise a good part of the vision.

Also, although there is a fundamental need to complete the technology work in parallel with the business work, it is clear that departments are already deploying solutions and will likely continue to do so at an increasing rate. There is, therefore, merit in identifying some early components that can be rapidly deployed (provided that there are appropriate standards in place), reducing the rework that will be required later.

The frequency of build of components (here, components are significant systems rather than individual sections of programme code within a system) in our EntArch may look like this:

How many of?	How often?	Example Components
1x	Very few	Authentication engine <i>Web services broker</i> Citizen data locker (or vault) Government directory (for people) Internet search engine

How many of?	How often?	Example Components
A few of	Many times	Rules engine Forms engine Content management system XML translation engine Circumstances engine Backbone integration layer Notification engine (for email and SMS), Records management system Intranet search engine CRM application DTV content publishing
Many of	Few times	Backbone integration adapters and transformation tools

There are really two key points with this table: (i) that the idea of building everything one time only and deploying across every piece of government is far-fetched; at the same it draws attention to the fact that this is indeed possible for a very few systems for which a credible business case can be made to be built in such a fashion. (ii) The default proposition when considering building something should not be to build from scratch or to buy (yet another) package and configure from scratch, but to take one that has already been deployed and reuse that instance – so we will have a few components that are deployed many times in different parts of government.

Perhaps the most important of these components, or at least of the unbuilt ones, is the web services broker. Although web services are relatively new, they are being deployed with increasing frequency in the private sector and many departments are actively considering implementation or moving to pilot. There are, however, two main issues: (i) the standards for how to access them and share data with them are not yet fully ratified, meaning that progressing too early means that the implementations may need to be changed later and (ii) the process for third parties (whether they are other government websites, such as the OGS, or intermediary services) to discover which services are available and then gain secure access to those (secure in the sense that government knows who is access the service and secure in the sense that the data exchanged is appropriate) is little understood. This is not dissimilar to the reason that the government gateway was originally built- it was felt at the time that if intermediaries were going to take up services, they would need to know how to get to the services concerned and how to exchange data with them. The gateway already sees traffic from a variety of intermediaries and more are expected to connect over the coming months. If web services proliferate in the same way that websites have to date, the risks are that few will be able to find and connect to the services, security will not be adequately deployed and costs will be correspondingly higher. Allied closely to the broker will be a set of management tools that understand how to react when any given service is not there, how to handle capacity peaks and so on.

Why Hasn't It Worked Before?

In attempting the delivery of a comprehensive Enterprise Architecture, we will face many new issues but also some old ones. Many people reading this document will be wondering how, if it hasn't been possible to deliver an EntArch at a departmental level, we can even dream of raising the stakes and commissioning a programme to deliver one for all of government. Likewise, the press will note that government's track record in IT project delivery is poor and that combining a wide-ranging business change initiative with a massive IT programme is even riskier.

There are some key guiding principles that will significantly affect our ability to do this programme:

- **Ambition.** *We're going to have to want to do this.* There are few examples of joined up projects in government and even fewer examples of successful ones. The ambition will have to be conjured up from the very top of the organisation, both political and administrative, and given the highest priority.
- **Choice.** *Choice is bad.* Giving departments the choice to opt in or out of this programme for individual projects or whole programmes will undermine it from the beginning. The might of the administration will shift from getting it done to finding ways to preserve the silo.
- **Business.** Technology projects led by technology people result in deliveries that do not maximise business impact. A new breed of people will be required to provide the *business leadership* that is necessary to deliver on this scale.
- **Customer focus.** Policy silos inside departments can and do create conflicting requirements that place enormous burden on those trying to comply. The overlapping windows of PAYE, VAT, Corporation Tax and Self Assessment make running a small business painful. The consequences of such windows can also impact other government products – the tax credits reconciliation process is made far more complex by a retrospective look at earnings, for instance. Undeliverable policy is a waste. Policy that does not exploit what is there is a waste too. Policy after delivery is no good unless it reinforces what has already been delivered. But without good policies, the initiative will fail. *Policy must be turned to focus on the customer consequence* and away from the silo.
- **Engagement.** Vendor partners have always followed the silo of government, matching their organisation to ours. *Vendors will need to be engaged* in a new way so that their efforts align with the programme. There are a few dozen vendors versus a few hundred departments, so there is a 1 to many ratio here that will give a big payoff.
- **Re-use.** Almost nowhere have systems or processes developed in one function been used in another (and this can often be the case *even* inside a department's multiple product streams). The concept of re-use, at both a component (individual system) and even a sub-component level preceded by a

rationalisation of the requirements process will ensure that “*unplug and replace*” is viable – reducing risk, time to market and cost.

- Asset building. Instead of building systems for individual products (e.g. PAYE, VAT, Council Tax) we must move to building assets that can be exploited by many products. Allied with the principle of re-use, this is the single biggest shift that we must engineer – the *move away from project delivery to programme delivery*.

Information technology projects have failed in the past because principles such as these were not adopted within a department or across departments. Not only have we often failed “to do things right”, but we have failed “to do the right things” – our context was wrong.

In the past, it has often been the “schedule” that drove how a technology project was implemented – a commitment had been made (often at the most senior level) to deliver something by a certain date, perhaps to align with new legislation. That commitment means that the heroes of a project are the ones that can most often deliver to the schedule – sacrificing scope, interoperability, performance or whatever. Delivering in such an environment is indeed a talent, but it will not advance an Enterprise Architecture. The new heroes must be the ones that can maintain the context of a project, recognising that exploiting assets that already exist will reduce the risk of failure, perhaps occasionally sacrificing the schedule. The consequence of the EntArch is that early on it is highly likely that the schedule will suffer most often, until there are sufficient components in place to exploit. This consequence must be knowingly undertaken and allowed for – if it is not, then efforts to build components will be hampered by departments seizing the chance to continue their own, individual delivery agenda without considering the context.

Ultimately what is trying to be addressed here that is different from what has gone before is a few big design issues. Creating a set of inter-operating components and the associated standards that *must* be used will advance the agenda significantly, whilst buying time for departments to re-engineer the endlessly complicated backend systems. Departments will need to show though that they are ready for taking this step – a sort of “capability maturity model” assessment. OGC’s Gate reviews, coupled with the Centre of Excellence initiative will be the ideal place to assess this.

What Are The Issues and What Should We Do About Them?

Despite the principles above, there will continue to be issues, some allied to non-observance of the principles and some as a consequence of them. At every step, close attention must be given to mitigating these points:

1. Technology before business. The EntArch is supposed to flow from the business strategy through the business processes into technology architecture and accompanying data structures and systems. That presupposes that there is an interlinked business strategy across the organisation. Despite some early efforts, it’s unlikely that such a full strategy will emerge in the timeframe that technology

investment decisions will be made for the next wave of service delivery. Some pragmatic decisions are going to be needed to ensure future interoperability and to reduce wastage, whilst still allowing the disparate businesses to proceed along their implementation paths – now is not the time for planning blight.

Laying down some core disciplines about what is and isn't allowed now will pay dividends later as rework is reduced: enterprise strategies and their resulting architectures in individual departments can be designed in at least broadly common ways. Design is, by necessity, a process of making compromises. Too often, an initiative such as this will get bogged down in finding the perfect way to handle a given problem – with the consequence that delivery has moved on and perfection is no longer an option. Settle for something that is good enough first and then allow subsequent iterations to improve on it. Provided the principle of “unplug and replace” is being adhered to, there will be opportunities later to improve the position.

There are many consultancies that will assist with the development and/or delivery of an EntArch – using various well documented frameworks as their basis. Any one of these could help avoid some of the obvious pitfalls, although it's unlikely any of them have taken the leap into something as complicated as the UK government.

2. Governance. It was noted earlier that the defining characteristic of the EntArch programme must be charismatic and capable governance. Poor governance will kill the programme, but only after much time, money and energy has been wasted. The governance model for this programme must surround it fully, providing an effective body to get decisions made quickly and with sufficient power to get the results of those decisions pushed through across the über-organisation. The governance arrangements must consider all of the stakeholders yet not get too big and unwieldy as to prevent progress being made quickly. Existing programme boards, champions meetings and so on have yet to prove that they can truly *drive* progress, although recent efforts at the eGDP (particularly with the Online Government Store) have proven that it can be done.

In reality, a network of governance arrangements will be needed:-

- At a senior business level, decision makers will review the business processes that are to be targeted, assign the resource needed, review progress and ensure that issues encountered in any single workstream are quickly addressed in others. Such a team will likely be headed by “Transformation Chief” who has ultimate accountability for the change programme. He or she will have pay review rights over senior business people across government.
- At a senior technology level, a pan-government CIO will be needed, with dotted line responsibility for departmental CIOs and again input to the pay and bonus process. The over-arching CIO will facilitate the rapid

conversations that will be needed to identify standards and push them through the organisation. The CIO will also put in place programme arrangements to ensure that projects launched conform to the technology standards, that innovation is encouraged but harnessed for the greater good and that the inevitable pilotitis of which all organisations are guilty does not cripple the programme. The aim will be to pilot where needed, but only where a route to scale is identified up front, with criteria that will define when and how to make the most of what is there.

- At a senior commercial level, the OGC will drive new processes through the organisations participating to ensure that harmonious buying processes are followed that allow any one organisation to take advantage of what any other has already done – this will include blanket licences, framework agreements and generic OJEU procedures for specific functionality.
 - Throughout the departments, various technology and business representatives will come together regularly to review common processes, opportunities to work through a consistent approach, lessons to learn and to understand and manage the ramifications of the change.
3. Requirements standardisation. Delivering one system for one department is a challenge, trying to deliver one for many departments is even harder. For this to work the requirements process will need rationalising and departments will need to focus on the yield of any given piece of functionality (i.e. the tangible change in productivity or real citizen benefit that it gives) – no yield means no chance of it being delivered.

The work done on central infrastructure highlights the risks here – if departments develop their requirements in isolation and then compare against a central solution, the odds of a match are staggeringly low (even understanding the words used by each party is hard). The consequence of this is that departments fear that they have to compromise on what is needed to fit in with what has already been built and therefore are going to disappoint their business managers. That said, the Standish Group have done some excellent work that shows that (out of a survey of 30,000 projects) only 7% of features built into a product are always used and a startling 45% are never used ... so if we could just get the 7% out of the way in the first release, we would instantly start to deliver a return on investment.

Changing the natural tendency to work in a silo is closely allied to the overall business strategy issue noted above. Unless there is a consistent application of business need across the über-organisation, aligning requirements will not be simple. The most likely solution to this problem is to pick certain (and very fundamental) horizontal processes that cross several organisations and use those as the base processes to start with – delivering rapid requirements for those. Early candidates might be accepting payments, processing forms (for validity), determining eligibility (for benefit), managing a claims process (e.g. appeals or

arbitration). Some of the opportunities that this might create are covered in the appendices.

4. Collaborative effort. Joined up government has so far stumbled at elementary steps despite good efforts by a few. A small number of departments collaborated to design the Government Gateway, many work to construct applications on the Knowledge Network. Critical mass though, a situation where departments willingly give up their home grown systems and donate their budget to the department charged with delivering the common solution, are unknown – Loch Ness monsters perhaps ... thought to exist but as yet not definitively sighted.

Departments note often that working as a team slows things down (you have to progress at the pace of the slowest and so on) but, by contrast, working alone means you don't get the benefit of wider experience and are more likely to make mistakes. For the key processes noted above, small (and very focused) teams should be brought together with a variety of technical and business experience to identify what needs to be done. These teams will need resources back at home office made available to them so that they can quickly find missing information and get questions answered. Participating in such a team (indeed, being a part of the whole EntArch initiative) should be seen as a badge of honour – these people will be helping to shape the future of how government works.

5. Communication. First, what might this initiative be called? “EntArch” whilst simple to type is not very catchy. Other terms such as “re-engineering” are a little 80's now and, if anything, will conjure up an effort to cut jobs. The name of the programme will not matter much to those working on it, but it will matter once the name becomes public.

More importantly, this is a major change initiative and will require a proper communication strategy that involves widespread conversations with staff, the unions, the business and commercial sectors and citizens. There is every chance in such a disparate and autonomous organisation as government for individuals to ignore the message, take a different path or delay action – whether through ignorance, a desire to retain control or an apparently conflicting directive. With the work on central infrastructure, we've seen that despite “best efforts” to put a consistent message out, the audience is too large and deep and the counter-intuitive gossip cycle too strong.

6. Funding. At the root of the lack of joining up is the absence of processes to joint fund and manage projects. Departments are restricted from moving funds freely between each other for a given project and if they are unable to demonstrate rigid value for money at each stage, are unlikely to want to commit. In an environment where Government IT project failures are plastered across the front pages most days of the week, it seems hard to understand the value for money argument – surely better value can be obtained by working together and using the small

number of great people each department has in any one field to deliver something that gives true value for money, both for the citizen and for government.

7. Centralisation versus Decentralisation. Many departments will see an EntArch initiative as an attempt to take control away for key components of their delivery agenda; or as something that will reduce their ability to deliver on Ministerial commitments.

Managing the shifts between centralised (top down) and decentralised (bottom up) projects will be challenging, but it is feasible. The biggest corporations have already faced down this dragon and realised that it can be done.

Remembering that the implementation of an EntArch is not a project but a long-term programme, the key to resolve this is to involve as many people from departments as possible in “flash” change programmes, where a sense of real ownership is fostered. Top level and very visible commitment will be required and rapid communication up and down. Governance, as always, is key. A highly capable “enterprise programme office” will need to be put in place to manage the interactions and dependencies. Such an office is in its infancy for e-government right now, so early lessons are likely available (as are lessons from corporations that have already pursued the EntArch strategy).

8. It’s not there. The last issue results from a difference in expectations. When central infrastructure is talked about there is some vague expectation that it will turn up in a box with a single CD and a couple of manuals; versus an expectation when procuring a solution from a vendor that there will be 6-12 months of contract negotiations, several months of design work and then a gradual implementation. The issue will be similar for “1x”, “few x” or “many x” initiatives in this context.

Departments who subscribe to the “it’s not there therefore I can’t use it” argument will proceed with procurements, technology delivery schedules and other programmes without reference to the core standards. That will delay the realisation of a wider EntArch for the sake of a departmental silo goal.

Although central infrastructure has matured in the 30 months since the government gateway was put live, it’s still what it is – an environment that allows certain things to be done well, given some work by the department and some work by the centre. It’s a bit like Microsoft Word. When you open it up, at your fingertips you have a wonderful set of tools to facilitate the creative process ...but it doesn’t write a novel for you.

9. Procurement. Government’s lengthy and apparently cumbersome procurement process is often a convenient scapegoat for why things are done a certain way. Products delivered in one department as a part of the overall EntArch must be made available for re-use by other departments without them needing to go

through lengthy evaluation or value for money reviews – there is no chance for every department to become a supplier to every other department here. Because procurement can take a long time to complete, there is always a tendency to throw as many requirements as possible into the specification to ensure that nothing is left out – this practice is the one that single-handedly can derail the EntArch initiative (see issue 3). An appropriate vehicle or set of vehicles must be in place to ensure that components of the Enterprise Architecture are readily useable by any department.

The Roadmap To The Government-wide Enterprise Architecture

There is a tendency with any new project to either spend a long time in the “feasibility” stage or to rush ahead with implementation. The former tend to take so long to do anything that by the time it’s done, things have moved on; the latter tend to deliver great, but flawed, products that need to be reviewed and rewritten not too long after launch – or, worse, a new version is written alongside the first one, but the first one is never killed off.

With a programme of this scale, a careful balance needs to be struck. If technology implementations rush ahead before standards are agreed we will be in a difficult position. For those who doubt the importance of standards, you need only think back to the Mars Climate Orbiter probe in 1999 where two teams working on different parts of software caused the Orbiter to crash into the surface of the planet because one team was working in metric and the other in imperial measurements. When the time came, the difference between “1” when measured in feet or metres is particularly significant as you come into land. At the risk of labouring a point, another error, in this case the difference between expectations between a “send” programme and a “receive” programme caused the loss of the Ariane 5 rocket in 1996, built at a cost of £6 billion. Ironically, the backup programme designed to provide fail-safe capability made the same error – it was running the same code. At the time, the investigation board said “that software is an expression of a highly detailed design and does not fail in the same sense as a mechanical system” meaning that potential failures are more subtle and harder to find than perhaps they might be in a machine.

Similarly (and at a very trivial level) in government, if two departments create a joined up application and one uses a date format of “dd-mmm-yy” and one uses “dd-mm-yyy”, only confusion can reign. On top of that, as services proliferate, changes between versions need to be rigidly controlled to ensure that the impact on other users of exposed services is fully considered.

Work on clear standards, data definitions and plans for significant amounts of reuse of proven software will be necessary to ensure that data can be safely exchanged, that errors are reduced (through maximal testing effort on as few versions of the software code as possible) and that a single failure has as few consequences as possible on the wider EntArch eco-system.

There are, therefore, some important (and upfront) pieces of work that need to be kicked off alongside the heavy lifting of governance, measuring success, initial targets and funding criteria:

Standards

For several years, various departments have struggled to establish both data and process standards in their own organisations. Some central bodies have tried to layer on top of that work proposals for common standards dealing with, say, “name and address”. Imposing standards can be dangerous: too many standards and confusion results, too few and chaos results; the right amount causes a degree of creative tension in the organisation but allows progress to flow. Little progress has been made in government to date.

One effort that has met with some past success, e-GIF, has now reached the point of uselessness as it has attempted to encompass huge swathes of emerging technologies and yet deal with the need to be fair and even-handed in dealing with every vendor. e-GIF is now a catalogue of varying standards, some of which compete with each other, few of which interoperate and many of which have yet to be firmly ratified along with a set of views on interesting and upcoming technologies of which readers might want to be aware.

All of the work must be harnessed now, however, to define two key types of standard:

- **Outward facing.** Between 1894 and 1904, over 6,000 telephone companies were founded in the USA and the number of phone lines went from 285,000 to 3,317,000. Customers of one company were rarely able to talk to customers of another company: standards were absent, competition meant co-operation was lax and customers were insufficiently vocal to get change to happen.

It took government intervention in the form of “independent, intelligent, considerate, thorough and just” regulation to resolve this issue, through the creation of a virtual monopoly. AT&T connected other companies to its network, divested some other parts of its business and established “universal service” throughout the nation (although it took until 1969 for 90% of the country to have a telephone at home).

Such a monopoly is not required in the world of Government technology - whilst fleetingly attractive to the centralists amongst us, it would never succeed – but the standards that underpin one are vital.

These standards would govern outward facing transactions – i.e. any operation outside of the department – so that they are managed in a consistent way so that 3rd parties, applications and portals can interact more simply with

government. With systems being increasingly “externalised”, i.e. made available to people outside of the government office (think mobile workers, suppliers of goods via procurement and intermediaries as well as citizens and businesses seeking service), there is a vital need for these standards.

It doesn't, for the moment, matter what happens inside the department – transformation of any given data format to any other format can be achieved simply and on the fly. This piece of work means that a standard address for all online services to be used must be put in place, a standard for accepting post codes, a standard for reference numbers etc.

In the future, much of this work will be wrapped up in the CIP programme and it may therefore make sense for it to be managed from within that programme early on.

Through this e-GIF will become a concentrated piece of work on schema standards. Other work, covering emerging technology can be put into an “Industry Watch” brief to keep people aware of changes that may be needed in the future.

For clarity, this is a piece of work about “data schemas”, not the data itself. Data belongs to the owning department – although it too will have to be dealt with to make the EntArch successful. This will be covered later.

- Interoperability – there are many emerging standards for technologies such as web services, few of which are ready for widespread deployment although many departments are already progressing to pilot implementation of products using them.

We will need to put in place an R&D function that will evaluate such issues and opportunities and issue “real standards” that can be deployed by departments that we can guarantee will allow interoperability, along with managing upgrade paths as versions change and functionality is enhanced.

These standards will ensure that differing services from within the über-organisation are readily usable by outside organisations (as well as by other government departments seeking to join up services or exchange information, such as the Online Government Store).

e-GIF should continue to be the repository for standards under this heading, making it a one-stop repository for two key things: (i) what format to use to get data into a departmental system and understand the response and (ii) how to communicate with the systems that government is using and manage the flow of transactions to ensure reliable delivery.

Data

Most conversations about technology (or, for that matter, new business initiatives) in government soon come down to a discussion about data quality issues: how many duplicates are there, is the data of high quality, what are the plans for cleaning it up and so on. As we advance in our initiatives to put services online, the risk of our “dirty data” being made available for citizen access increases, as does the risk of exposure. It’s hard, in fact, to see how to progress efforts towards and EntArch if the underlying data is of poor quality – if we offer pre-population services and display incorrect data, or joined up services but are unable to match up a person A in department X with person A in department Y, the issues will quickly be exposed – at least for that individual.

Tracking back, there appear to have been many initiatives to clean up data. For the large departments, such a programme would be enormous – not dissimilar perhaps to their Year 2000 efforts and its therefore understandable why few have succeeded (without a looming deadline or external event to focus thought on such an issue, it would be easy to defer it for another day). The CIP programme will eventually tackle this too, at least through its provision of a common identifier for citizens (there are, as yet, no solid plans for a common business identifier although ONS have published papers on this in the past which may be taken up, perhaps by the business.gov initiative).

One of the main issues with data is that it rusts – the less that it is used, the less likely it is to be correct. Some data items remain impervious – your date of birth is unlikely to change (subject to keying errors and ignoring the rumour of a difference between DWP’s benefit calculation and IR’s tax calculation for recent immigrants that requires different dates to be entered), but items such as your address, your GP’s name, your employer and so on may change often. Indeed, even your name may change, through marriage or via Deed Poll.

Dealing with such data rust could be dealt with by giving the citizen/business access to their data in one place only, where they populate the initial loading. Subsequent interactions are always compared with the data stored in the citizen’s “locker” and differences exposed for correction. The standards referred to above for correcting variances in format will continue to be necessary for some time to come. The citizen data locker will be considered later.

A second big issue with data is its rate of growth (various research bodies have forecast growth rates of anywhere from 75% to 100% per annum for large organisations – so a department with three terabytes today [not an uncommon amount] will need to deal with 300 terabytes in a few years time. Storage hardware will undoubtedly grow to meet that need, but storage management capability has historically trailed.

To stop the problem getting worse, departments are likely to conduct their own cleansing exercises on various databases and, at the same time, should institute

data standards to facilitate access by future initiatives. Trying to come up with a single standard across government is an unlikely endeavour and so the focus should be on known standards in each department that can be readily interoperated with – a series of “Enterprise Data Architectures” or “data strategies” that will help rather than hinder the wider initiative.

A third issue is one of “data alignment” – the differing windows used by departments to calculate taxes, tax credits and benefits. For example, VAT is quarterly in arrears, PAYE is monthly in arrears, Self Assessment can involve two payments on account with an annual reconciliation, Corporation Tax is anything from 8 to 20 months in arrears, depending on financial year ends. To successfully align business processes, much thought will need to be given to how to reduce the variety of such misalignments, simplifying the engagement process for the citizen and the business and reducing the management overhead of differing data sets.

The single biggest programme to impact data over the coming years will be the Citizen Information Project (CIP). The plan is for a single database to hold a Personal Public Service Number that initially cross-references to all other numbers (such as NHS number, National Insurance Number, Passport Number etc) but that ultimately becomes the key reference both inside and outside of government, perhaps tied in to identity cards. The project has just exited initial phase and will be in the “definition phase” for the next 15-18 months, emerging with a procurement in early 2005 and implementation from 2006 onwards. The key data sources in stage one will be the Passport Agency and DVO with Inland Revenue and DWP following later, perhaps in 2007/2008. This is an enormous project that will shape the future of the EntArch (for citizen facing services). For CIP to clean up the data itself though (and then to construct a mechanism for passing the cleaned up data back to the source department and then yet another process for handling mutual updates of future channels) is a significant undertaking that may be beyond feasibility for a single project unless action is taken to prepare the way for the project first.

Importantly then, CIP will need many data standards to be in place, it will need integration backbones and it will need transformation capability – all of these streams of work defined in this EntArch. So there is much work outlined here that will help CIP deliver faster than it might otherwise, and there is much benefit in the CIP team being a key part of the EntArch team to make sure that the work is interlinked and that no disparate decisions are made.

The Council of Architects

There are few occasions when architects from around government get together to discuss and agree implementation of solutions. Any one department can roll out an entire new service that may be innovative or (worse) may be fraught with issues, without any other department being aware. If it turns out to be a great service, few can take advantage; if it doesn't, few find out what went wrong and avoid the same mistakes. OGC's Gateway Reviews, whilst being of enormous

value for individual projects, do not operate at a technology implementation level, or at a technology programme level (are they doing things right and are they doing the right things?) and do not provide the information on issues that they do find to a wider audience. Changing this approach will be a big step.

Major projects for departments need to be reviewed (in hit squad fashion) by a council of peers – people who are intimately aware of the potential for benefit, the issues that might arise and the opportunities to learn lessons. These reviews may be brutal in that they cause delays to projects that are seen as off track, or they may be quick – but in either case, the lessons learnt, the techniques used, the opportunities to copy and complete will be made available to the rest of the council for reuse.

Peer pressure will encourage the architects present to reuse techniques that have been proven elsewhere rather than try and use obscure and unproven new ones. Products that are more hyper than substance will be found out in these meetings and architects can work out solutions that may lead to longer term synergies.

Likewise, there is increasing recognition that it is better to buy (reuse) rather than build each time. It was long ago realised in the motor industry that the production line with significant component reuse wins over bespoke, hand-crafted solutions (if you're buying a Porsche Boxster, you're getting 60%+ of a Porsche 911 but paying only half the price – meanwhile Porsche is making margins of 50% on each). The problem, of course, is that if every department buys different packages then there may be little opportunity for reuse. A strategic review of package solution usage would need to be accompanied by a comprehensive review of business process in many cases – e.g. aligning financial reporting around a single version of Oracle Financials would require dramatic changes in every financial department not using it (and, for that matter, likely in the ones that are using it as the odds of a common set of features being available is low).

Next Steps

Once the early decisions have been made clear, some of the detailed work can take place. Four early streams have merit:

- Putting in place the organisational competencies to deliver the programme
 1. The Programme Management Office. It's very likely the EntArch PMO can build from the OGC's Centre of Excellence initiative and the work that OGC already do on Gate reviews. There will need to be a new focus on business process in the widest possible context as well as the rights and wrongs of technology, but given the progress made so far that should not be difficult to achieve. As noted above, project successes have to go from being seen as heroic to being systemic.

2. The Technology Policy team. There will be many policies that underpin this initiative, some will come from the centre and some will come from departments. It would be foolish to subject all of these policies to lengthy consultation and review – some of them are just going to have to be decided and made policy early on, to prevent departments from implementing standards that will be incompatible with other departments. The point of components such as the integration backbone is to make the backend technologies as immune as possible from the front end interactions, meaning that departments can work their own data standards and clean up projects without reference to the outside world at least early on but that, at some point, the new components built will have to conform to central data standards.

Technology policy should therefore be a small team in the centre with a reach far into departments to quickly develop and ratify standards. Industry squabbles over the merits of given technology should be ignored and a focus put on what is available to be used now. If it's available and it fits the EntArch strategy, it should be endorsed and we move on.

- Fleshing out the standards process and the standards required
 1. Schema Transformation. Each backend operates differently now. Departments have spent a lot of energy to date building translation tools that take XML (the new and close to universal descriptive language used for Internet applications) into back end-speak – but each department has built their own with much hand-crafting. To make real progress here we need a stream of work on “integration backbone” systems – a few systems that can take any input and map them to any output.

Some departments will want to be able to do that in absolute real time (synchronous) and some will want to do it in batch mode (asynchronous). The technology already exists for the latter in the government gateway but is little used by departments. Even with, suppose, a single integration backbone the biggest amount of work is in the “adapters” – the pieces of code that understand the XML for each process (each online form, each interactive query etc) and translate that to the specific backend in use for that transaction in the department. Many of the necessary schemas are not defined yet, and there is some risk in running ahead with defining such schema when the shape of the business process after change is not known. Still, the risk of planning blight is greater and moving ahead here will allow some early deliverables.

The key point is that every department should not have to build such a backbone – a few will, and their work should be cloned across the remainder of government.

2. **Data.** This is probably the hardest problem to solve. To date the problem has largely been hidden from the mainstream as there are too few services available online. That said, online registrations for the government gateway show at least a 20% failure rate for tax transactions because of out of date addresses held by the departments involved. For an annual transaction that is perhaps to be expected, but it still shows an alarming amount of data rust.

Each department is doubtless already reviewing their plans for data. What is needed is a pan-government initiative to review and reduce the “Total cost of Data ownership” – that’s not something that is talked about much, after all, it’s far easier to talk about the cost of a desktop or the cost of a system. The data is the real asset, and the widespread duplication, inconsistency and just plain old data has made the cost of managing it exorbitant.

A rapid-fire working group set up to address the problem on, initially, the few chosen business process schemes needs to be set up. Its terms of reference need to be clear: there’s no such thing as “too hard”, the goal is to reduce the cost of ownership of our data by 25% (or maybe even 50%) and to deliver a sustainable plan that will keep the cost low as volumes grow.

3. **Interoperability.** With the work on data quality and transformation between services underway, the next problem to address interoperability. That is, how do different services work together in ways that today might seem far fetched but that in the future should become the norm. For instance, how would we create a transaction that allowed you to purchase your car insurance, tax disc and parking permit for outside of your house in a single page; or how would we create a transaction that allowed you to register the birth of a child, apply for child benefit and child tax credit all in one go?

Probably the most important standard for achieving this in a world where we are working to an EntArch are the web services standards: which services are available, where are they, what parameters do they need and how do I talk to them? If we don’t get those right (or very close to it) in the beginning, then any attempt to join up disparate services in the future will fail dismally unless the offending services are pushed through large amounts of rework.

There is, therefore, a strong case for a single “web services broker” (the technology entity that addresses many of those questions) underpinned by a set of government standards (which may, initially at least, be different from industry standards because those standards are still undergoing changes and are far from mature). Such a broker could have a single

standard to the outside world that conformed to whatever the majority view of a mature interface was, while handling various states of government standards as different organisations reach different levels of maturity.

- Mapping out and modelling the business processes as they are
 1. What's the business flow?
 2. How does technology support that flow now?
 3. What are the overlaps across the initial processes being reviewed?
What are the opportunities to collapse processes?

There are well proven modelling techniques that will allow rapid progress to be made in identifying the stages in each of the main processes chosen for first look. Some departments, notably HM Customs, have already made great progress in looking at how their business model looks, how it will look and what needs to be done to the technology to achieve that change.

- Agreeing a rapid requirements process
 1. To progress this difficult task, government teams will need to operate under a new kind of constitution – a set of agreements that all departments sign up to that helps big departments get what they need fast without upsetting smaller departments who might be more fleet of foot but have less money. It will include such topics are:
 - How will the process to collate requirements be managed?
 - Who are the key stakeholders?
 - What's required from each stakeholder?
 - How will competing requirements be prioritised?
 - What are the escalation routes when an issue remains open?

Once these are kicked off and producing output, some generic technical work can take place. This can happen alongside as, for the most part, it's likely that these components will be needed by one or more of the streams, so the earlier that they start, the earlier some of the issues in design will be understood.

- Agreeing the 1x components
 1. Who will own them? The e-Delivery Team is strongly position to own some of the them, but other departments may have services already

available that are suited to being exploited by wider government (for instance, the Camden/Local Government APLAWS project)

2. What are the requirements? OeE are presently reviewing the governance process for central infrastructure, a task that will also result in some proposals for how requirements can be built (that can usefully feed into the requirements process noted above).
 3. What will the procurement process be? It would make sense for a menu of procurement options to be available for building out new components – these should include procurement through existing supplier relationships, framework deals and more significant procurement exercises through OJEU. The work that Richard Granger has done in the NHS and Jo Wright in CJIT can be used as an effective model for the work that must be done more widely across government.
- Agreeing a “cloning” process
 1. For the “few times” components, who will own and clone the key components, making sure that there is effective reuse?
 2. What is the change control and release management process for cloned components? Once a component is deployed in the wild, it will be open for change – changes that might affect the ability to carry out further upgrades in the future. Where components need to be deployed in multiple places, there must be consideration for the security of their deployment, the update process (both patching and full new releases) and so on.

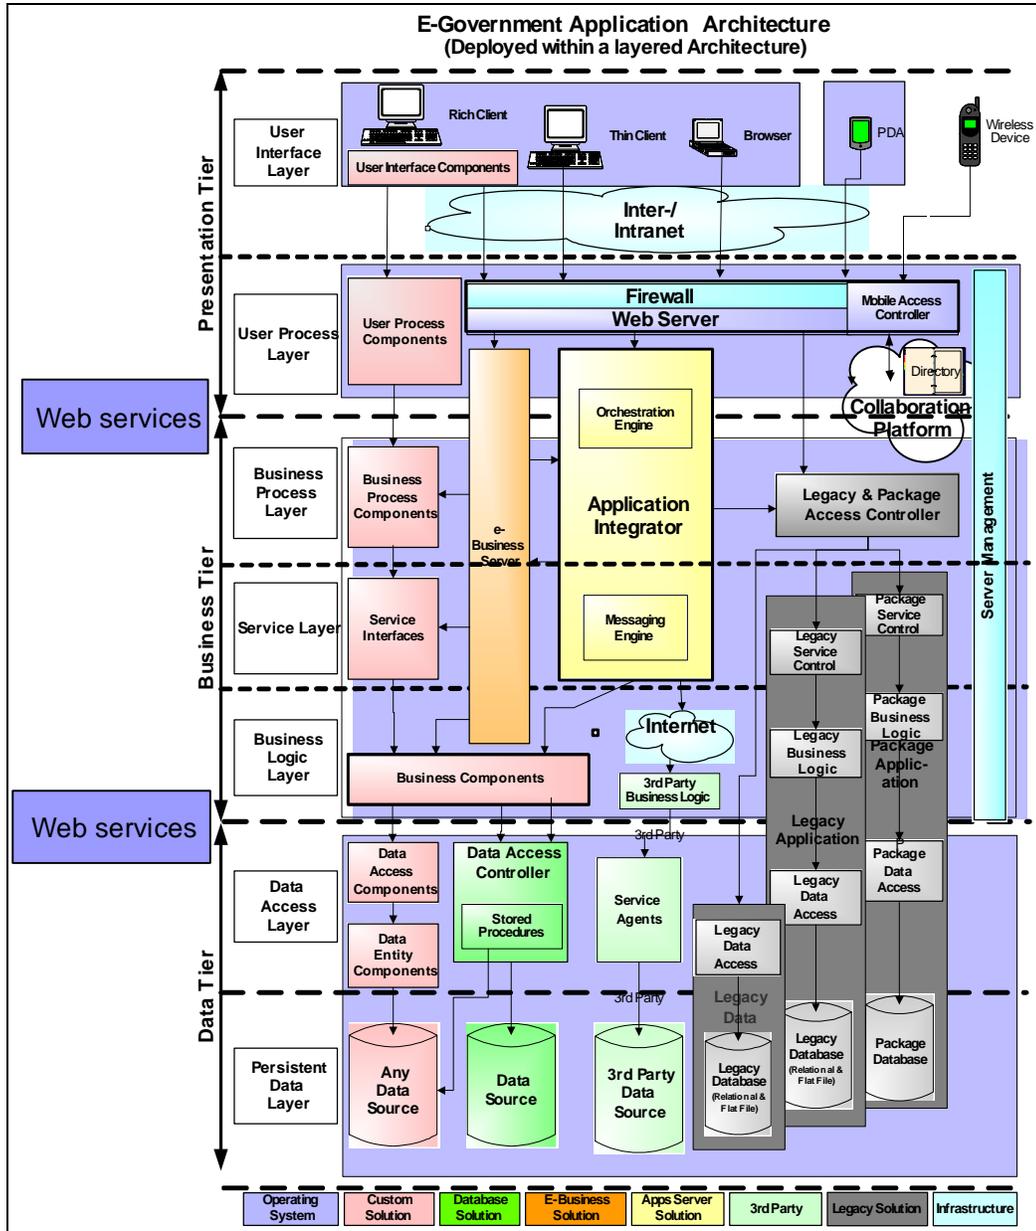
Whilst all of the setup work is going on, other teams must be involved in the next stage of work:

- What will the roll out plan be?
- What impact will there be on existing IT relationships?
- Which projects may have to be cancelled to make room for the new projects?
- What are the imperatives for departments that must not be impacted because of existing commitments?
- What are the processes that will be teed up for review in the next stage?

And so on.

Appendix A

A more complicated view of an Enterprise (Technology) Architecture

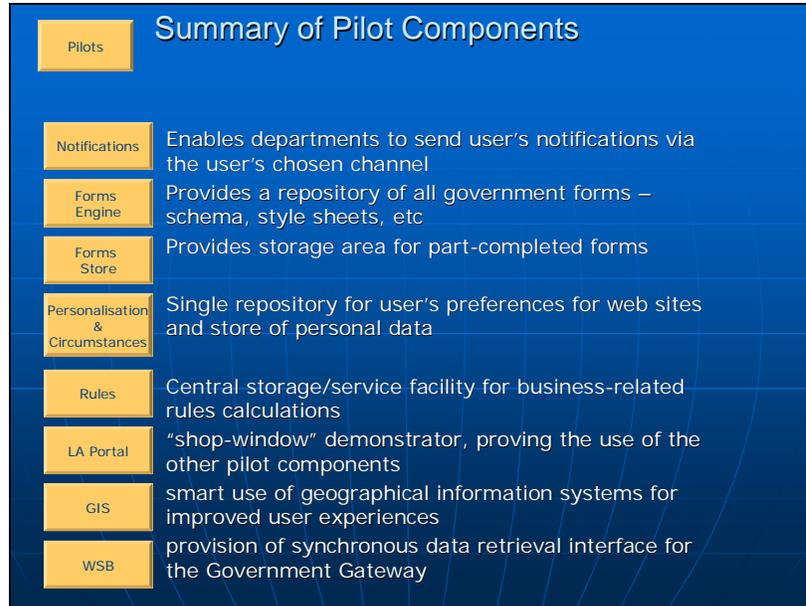


Source: Jerry Fishenden

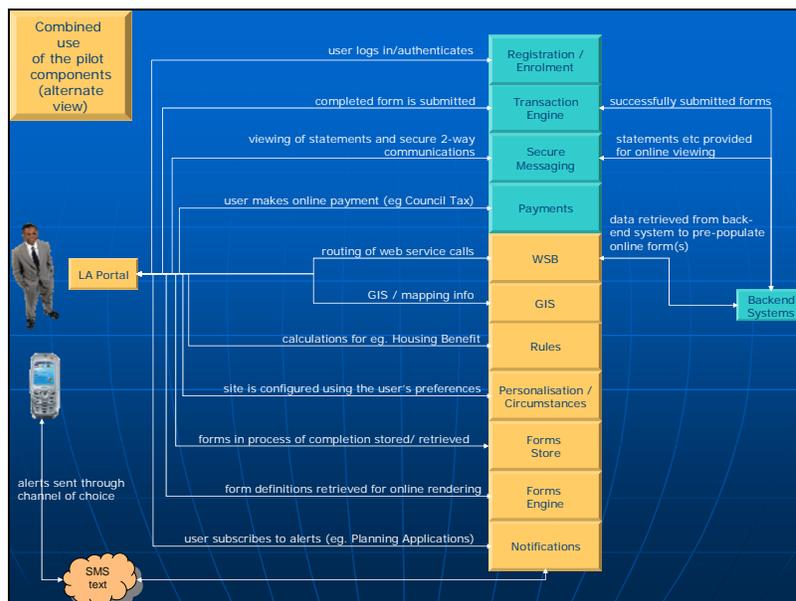
Appendix B

Common Components

In early 2003, the e-Delivery team made a successful bid for ISB funding to develop pilot components that could later be made available to the rest of government. The components identified at that time were:



It was envisaged that these components would interoperate with existing central infrastructure components, as in the following example:



The components identified above do indeed form some of the potential “few x” systems identified in this paper. Some consideration should be given however to a strategy allowing the following lines as the fastest route to market:

- Web services standards

Web services are set to become the new “must have” toy. Government has spent the last five years building thousands of websites, and these will quickly follow in that path.

“Web services” is probably a confusing misnomer for many. In absolute terms they have nothing to do with the web (new technology) but an awful lot to do with legacy systems (old technology). Web services are an evolution of prior attempts to “wrap” aging systems in a new layer of software that makes them conveniently available to the outside world – a process known as “exposing” the service. What this means is that new applications can be built that will talk to a web service rather than a legacy application to get the answer that they need for a given question; later, the legacy application can be replaced (either fully or component by component) and the new software placed behind the same web service, avoiding the need to make any changes to the new applications. This work, correctly defined and handled, has enormous potential for government. It can allow the front end to be comprehensively redefined to create a view of joined up government whilst buying time for the destruction and replacement, piece by piece, of the legacy systems that in many cases are more than twenty years old. This is new technology and not necessarily ready for full deployment on a scale basis, but the software industry is putting enormous backing behind the effort.

To ensure later interoperability (even if there are few joined up services now), clear *government* standards will have to be in place along with a watching brief on how industry standards are emerging so that the standards can be harmonized at the earliest opportunity.

- Web services broker

The broker is a node that acts as a repository, yellow pages, or clearing house for software interfaces that are published by departments. Ideally, there should probably be only one of these – competing yellow pages may cause confusion and present issues around ensuring that they are all up to date and available.

New applications, whether connected to the new technology in the web world or imitating the legacy world, will increasingly be built using web services. As the web services multiply, management of them will become an enormous problem. Applications will be built in the outside world that assumes these services are highly available, are secure and that they always return the right answer to a question asked. At the same time, malicious programmers will look for ways to

exploit weaknesses in the services. A massive effort will be required to implement web services in a clean, managed way with appropriate standards. Such an effort can only be done on a pan-government basis, else the chain will be let down by the weakest link and the reputation of the whole will suffer. It is the web service broker that will drive this work.

- Rules engine standards

So far, almost every department has created a separate processing infrastructure for each online service that they have put in place. That has led to multiple portals, multiple rules engines and a disjointed customer experience (visit PAYE and VAT for a view across two departments or Tax Credits and Self Assessment for a view inside one). This shows a lack of enterprise architectural planning within even single departments – that’s understandable in many ways, this is complex stuff that needs rigid and long-term planning. Most of the online services built to date have been experimental versions to test the water. Take-up has been low and these environments have not been stressed. Before we get much further, we need to rationalise the architectures in each department to reduce the cost of maintenance and update.

One of the key requirements for any online service is a rules engine. This is best explained by imaging an online form – one for self assessment for instance.

The form is constructed of a number of fields, each of which can accept various kinds of input. Some of the fields are purely numeric (e.g. salary), some are alphanumeric (e.g. address) and some might be binary (a yes/no field, like “are you married?”). Each of these fields has rules that can be applied to them individually – for instance, checking that a date is indeed a real date and, if it’s a date of birth, that the customer hasn’t typed in today’s date by mistake. They also have rules that can be applied to a range of fields – for instance, the salary should be higher than the total amount of tax paid, the tax taken on a dividend must be less than the dividend itself. These are “contextual” rules that understand the entire form and what constitutes a valid form. Finally, there are external rules that might, for instance, validate a name and address (perhaps through cross-checking with another system) or check that a National Insurance Number offered indeed corresponds to the name on the form. Together, all of these rules help ensure that when the data arrives at the back end system, it is valid and can be processed without exceptions (comparing this to the to and fro process with paper where multiple errors are corrected over a period).

So a rules engine is required whenever a service is put online, even for very simple services. The problem is that most rules engines to date have been built old-style, i.e. closely coupled with the service that they support and using hard-coded rules. That means that whenever a change is made, there has to be extensive programming work, testing and overhead. It also means that third parties (such as application providers like Sage or intermediaries in the future)

who want to use the service and choose to build their own branded application or web service will have to write their own rules engine. Experience shows that this is not done as well as it is within government and so more errors result from third parties than from government systems. It also means that any change made to the form, and therefore the rules, has to be notified to a wide-spread audience – and means that costs multiply dramatically as the consequence of any change.

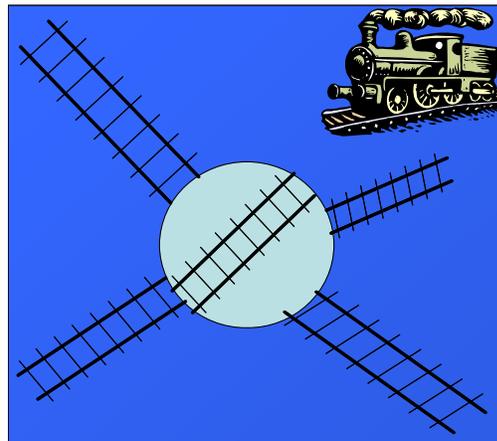
A few rules engines built and deployed across government and made available as web services could solve this. Each one would have the rules to sets of services, made available in “English language” so that they could be changed easily. Third parties and government could call those engines from their own online services and ensure that input was valid. Local authorities could more rapidly deploy, say, housing benefit transactions because after the first set of rules was defined there would likely be only small variations from authority to authority. The cost of maintenance falls, the cost of roll-out falls and the chances of valid input from third parties rises.

- XML parser / Integration Backbone layer

This component is often today coupled with the rules engine, for good reasons. It reads the XML that is generated from an online service and checks that it is valid (following the rules) and then, where necessary, changes the XML into back end speak so that it can be consumed by the legacy systems. There are as many of these in government as there are rules engines – pretty much one for every online service (give or take a few).

In the world of an EntArch, there will ideally be a single Integration Backbone for any given department, with the biggest departments taking the lead perhaps and making their systems available for wider use by smaller departments. Delivering such a backbone is a big task and one that should not be undertaken hundreds of times across government.

To try and paint a picture of how an integration backbone works, let’s try this example. Imagine a railway turntable where trains arrive from different directions, but that each track has a different gauge. The turntable lifts the body of the train from its existing gauge (leaving the wheels) and moves it to a set of wheels waiting at the next gauge. Every time a train arrives at the turntable, this approach is repeated.



To put this matter in real terms, the gauge of trains in Ireland today is 1600mm, in the UK it's 1435mm. Were we to create the equivalent of a Channel Tunnel to link the mainland and Ireland, we would need to lay new track at one end or the other to allow a train to pass through, or create the turntable-idea that I note above. Both are probably impractical and although standards on railway gauges were imposed in 1846, it was by then too late to solve the problem as the railways were already built.

In technical terms, we're in the same place. Some standards do exist, but not enough. Imposing them now is far too late for any system already built that does not conform to the standard (and that means almost every system we have, even ones in the same department). So, what is needed is a device to do the heavy lifting and make the necessary "gauge changes" every time they are needed.

Ideally, there is a defined standard for incoming and outgoing messages (see the pages earlier covering standards), so there is only one change needed per query – from the outward standard to one of the various inward standards.

Sticking with the railway analogy for a little longer, when the Union Pacific and Central Pacific railroads were joined in 1869, the 2,000 mile journey from coast to coast that had previously taken up to four months was reduced to a mere six days. Standards and integration drive significant customer benefit!

There are, occasionally, compromise methods that can be used to forward the integration path – Britain's own railways had competing standards early on, the Brunel gauge of 7' ¼" and the Stephenson gauge of 4' 8 ½". The incompatibility was solved by adding a third rail to hundreds of mile of track to allow trains of either gauge to pass – a complicated but workable solution. Far better though, surely, to agree a standard and stick to it.

As the older systems retire, they are built using the new (outward facing) standard and the integration backbone does less work. But that will take time. Until then, the backbone is a vital part of putting services online, of offering joined up services that are transparent to the customer and of buying time to replace the legacy architectures.

- Forms engine / Forms standards

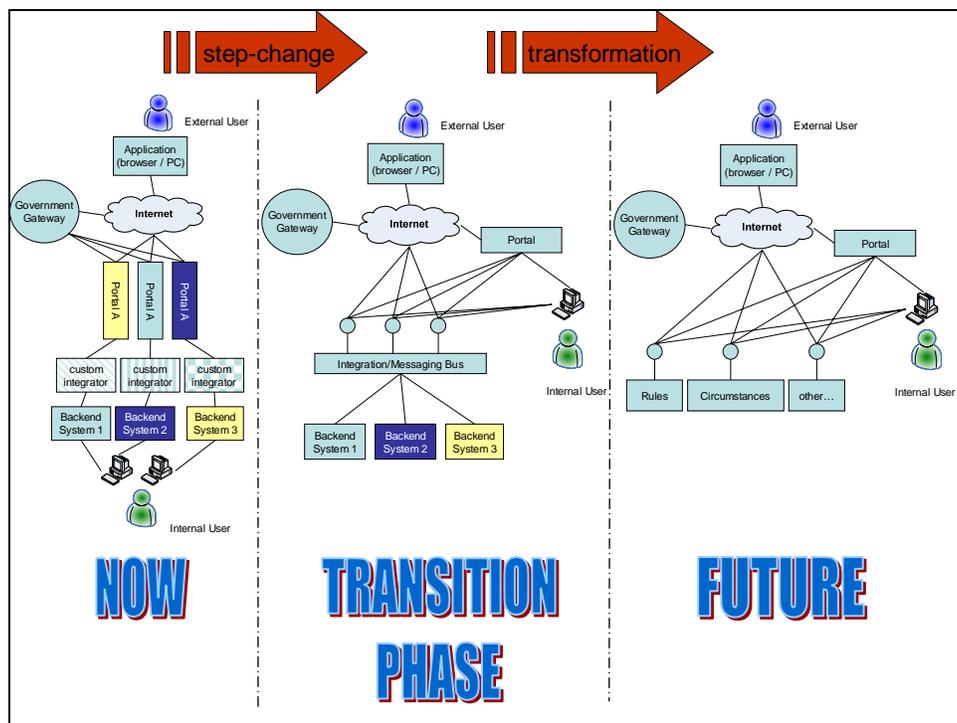
Forms engines do what it sounds like they do – they generate a form and display it on a web browser, allowing it to be completed by the citizen and then sent to government. As with rules engines, most forms are built by programmers making them difficult to change and therefore slow to respond to customer feedback. New technologies are emerging that allow forms to be constructed by non-technical people but, more importantly, allow forms to be stored as a set of parameters (with associated rules held by the rules engine) that can then be

displayed not only by the government department that owns the service but also by 3rd parties (such as the OGS or intermediaries).

It may be that an actual “forms engine” is not needed – there are many packages already in existence. What will be important is that they work in a consistent way and can send messages (the forms, wrapped up in XML) to government in the way that government wants – i.e. following the standards that we impose. Third party providers (such as Mandoforms, Jetforms, Adobe and so on) have already seen a big opportunity in government and are doubtless selling versions of their software that deliver the forms in some proprietary (although still XML-based) format – worse, it’s possible that each government entity has tinkered with the format provided so that it is not even standard to the vendor package. If the standards are developed correctly, then each third party has only to develop a single version of their software for government that can be more widely deployed without the need for tailoring for each entity. This should be a win for all sides.

Transitioning

The route to implementing these common components and/or their associated standards is complicated when viewed across the whole of government, but relatively straightforward if taking it a process or piece at a time. The diagram below shows how it might go, in much simplified form.



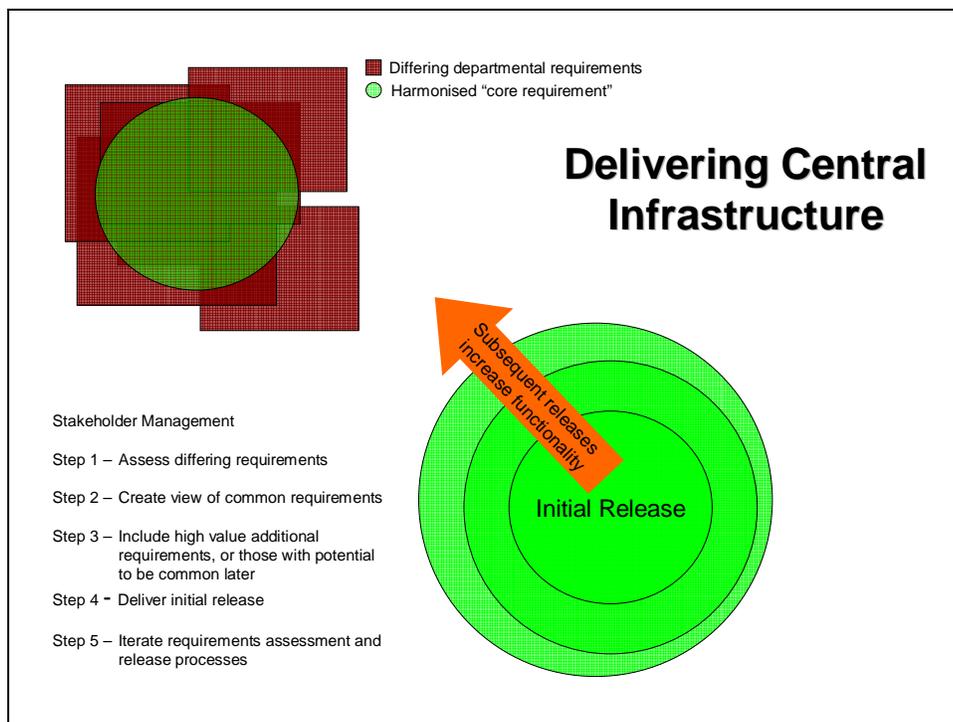
Source: Jerry Fishenden

Appendix B

Streamlining the Requirements Process

To make work on the EntArch process even a little bit feasible, the way that requirements are drafted will have to change. Experience with central infrastructure has shown that it is possible to draw a small number of departments to a common requirement set, provided that the development is being undertaken sufficiently far in advance of their thinking or procurement being started. If the comparison process takes place once a solution has been designed or a set of requirements has been drawn up from scratch then one of two things occurs, (i) changing the architecture adds risk, cost or time to the project or (ii) understanding the requirements, the choice of words, the gaps between what the department wants and what is available elsewhere effectively puts the entities in a competitive tender position where rigid negotiation positions are held.

The slide below was used to illustrate the dilemma facing builders of central infrastructure.



Sometimes the mismatches in requirements are the result of a genuine need for a given piece of functionality, sometimes it's the result of the obtuseness of the departmental representatives and, often, it's the result of a vendor position on a given piece of technology. All of these forces are present in varying degrees in each department that is approached to co-operate.

Appendix C

So, I buy into an EntArch. What will it get me?

Throughout this paper, I have largely concentrated on process, standards and a mechanism for ensuring that the EntArch sticks acknowledging the lessons learnt through delivery of central infrastructure and other IT projects in government. Here are some of the things that could be done if there was consistent, componentised technology architecture in government with widespread use of common standards – this is not technology leading business but technology enabling business.

One Vision

Today, the introduction of a major new policy initiative (Tax Credits, say, or Pension Credits) needs a new system to be built from the ground up. The system needs to understand the customer base (who can claim), the entitlement process (the various trigger points and the rules by which people claim) and then manage the payment process (coupled with reconciliation, fraud control, reclaiming overpaid funds etc). In the examples above, Tax Credits was a significant programme involving a huge new system that also linked to several other Inland Revenue systems; Pension Credits is being introduced through making substantial changes to the Income Support systems although a new system will later be commissioned.

With an Enterprise Architecture in place (i.e. one that is sufficiently widely deployed), much of this work (and the consequent cost, risk and lengthy delivery time) can be avoided. This statement is not made blithely – an EntArch is an enormous undertaking, but if it is commissioned correctly and delivered properly (in stages with the business leading), the benefits flow.

For instance, new credits can be introduced by adding rules to the existing rules engine. The rules would be relatively simple to change without the need to change software code. Changes to existing rules could be made in an equivalent fashion. If the new credit needed to talk to another system, it need only understand the web service to talk to – all of the translation between varying data standards would be carried out by the integration backbone. Payment would be handled through a single engine that carried out the reconciliation, managed recall of funds, understood duplicates and watched for signs of fraud.

Proposed Business Initiatives

To carry off something like the vision described above, some of the projects needed would be:

The Department of Entitlements

Entitlement is a well-used word in government. It covers every aspect of government from local through central and a range of dozens of different benefits, credits, grants and loans. Ultimately entitlement is a simple process consisting of many “If ... then ...” statements, i.e. “If you are between 16 and 65, then you could claim disability living allowance”, or “If you are over 65, then you could claim attendance allowance”.

The range of rules is complex and overlapping, but they are possible to describe in English and so can be adapted to a rules engine. The present paper forms for any one of these presently run to a dozen or more pages *each*, yet the amount of real (i.e. used to assess entitlement) information is small and is certainly largely consistent between the different forms. For every one thousand forms in government, perhaps there are 25-50 actual pieces of information that are needed, and most of these will be common.

Creating a “Department of Entitlement” would mean putting in place a single service where people could go, anonymous or otherwise, to see what they are due from government. The service would pose several nested questions that would lead to more questions, uncovering each of the benefits, grants and opportunities. The information would be stored, if the citizen wanted, and used to make the appropriate claim to the owning department. Each department would have exposed their claim process as a web service, so the “DoE” would not need to know what data format was used, how to talk to each legacy system or even whether the legacy system was actually available.

The navigation process through “DoE” would need to be given careful thought – someone applying for a student loan may not be eligible for DLA, so, where possible, existing information on the person could be drawn from other systems (again, the “DoE” need not know much about these systems, they will all be made available via web services) to reduce the number of apparently unrelated questions asked.

Is this possible? The USA have been operating an online services (www.govbenefits.gov) for over 18 months which goes some of the way. It asks the questions but is unable to complete the application process because the legacy systems are not yet connected via web services. But it’s only a matter of time.

Creating a “DoE” would require some significant policy shifts and a recognition that the consequence of it could be dramatically increased take up of benefits and credits, leading to an increased financial drain. That, though, is surely the point of having benefits and credits.

If the “DoE” were also to include in its role a “What’s in it for me?” section, it could expand to identify and explain other areas where government may be able

to help – for instance, how to seek government help where your pension has been lost through mismanagement, how to find out if there are any retrospective amounts owing to you or simply how to plan your finances more efficiently through tax efficient schemes, baby bonds and so on. Or, such additions could be left to the private sector who could seek to add value to the standard “DoE” offering by acting as an intermediary.

The “DoE” could operate just as effectively online as offline – asking such questions over the telephone, in a Citizen Advice Bureau or a Post Office, whilst perhaps a little more cumbersome than using an interactive online application, is likely to be a better experience for the citizen than trafficking from one government office to another over a period of weeks.

The Department of Who You Are

Government spends a significant amount of money assuring that the person asking for something is indeed that person. Identity fraud is a significant sum (put at around £1.3 billion in government reports) but it’s likely that the “prevention of identity fraud” budget across the whole of government is significantly more than that.

Today’s online services require separate identity checking at each department, using a range of numbers that represent how you are known by that entity – your national insurance number, your address, your tax reference number and so on. No single department trusts any other department to do that work. Although the identifiers can be joined up by the government gateway, the lack of cross-trust effectively means that every interaction is a new one. Joined up government doesn’t exist.

Also, no department has yet acknowledged the weakness of their own identification process and sought to forge an online linkage with 3rd parties who might have better ideas about who people are. Banks, for instance, have strict “know your customer” rules to help protect against money laundering. Credit companies conduct widespread checks to ensure that the risk of losing money that they lend is reduced.

With increased cross-trust, government would no longer be reliant on the lowest common denominator for identity checking (for instance, anyone who holds a passport can usually use that to pyramid up the rest of the layers of authentication, as evidenced in recent Daily Mail stories). Cross-trust means that a balance of probabilities can be established using a variety of sources which can only lead to a reduction of mistakes and successful frauds.

Once the CIP project gets underway, a single identifier becomes a reality, but being in possession of that identifier doesn’t mean that you are the rightful owner, unless it is somehow tied into fingerprint, retina or other biometric checking –

and, even if it is, it's the issuance process that then becomes the target (as it is today with passports).

This is not a trivial problem and the EntArch is probably only a part of the solution. What having an EntArch will allow is (a) greater cooperation between departments in checking identity through the ability to ask several questions of one person combining data from around shareable data from around government, (b) greater reliance on third parties (such as Experian, Equifax, BT URU and the major banks) through asking questions using their databases (such as prior address history) and (c) more dynamic questions based on current data rather than rusted data (so, questions like "what's your bank balance" rather than "what's your account number" become viable).

All of these can combine to give a greater degree of certainty that the person asking is who they say they are. The absence of a strong ability to do this is one of the great hindrances to the online programme, despite the efforts made in establishing the government gateway. Now, we need to exploit what we have using the context of the EntArch.

The other key piece of such a department would be the ability allow others to act on your behalf, until you say they cannot or for much shorter periods of time. Certainty over who the intermediaries are coupled with certainty over who the individual was who granted the right would help those on the wrong side of the digital boundary access services that could be highly beneficial – because, at all times, the aim of this EntArch is to offer the most valuable services to all comers through a variety of channels.

Such an authentication service could even become a revenue earner – the need to prove who someone is arises all the time. If the service operates effectively, it becomes the gold standard for which other businesses should pay for access.

The Single Account – The “Books of the Country”

As we componentise the existing technical architecture, removing legacy (monolithic) systems, we remove functions that are traditionally performed in a silo and make them pan-government. If we pursue that course to its natural end, the ultimate monolithic system that is left is the “books of the country” – the general ledger – that knows about all inflows and outflows of money and the balance at any given time, either for an individual, a business or, for that matter, the country as a whole.

Such an individual cash flow, or citizen account, would be similar to account aggregation as practiced by some banks today – the ability to see your current account, savings account, mortgage, credit card and stock investment plan in a single view. Incoming amounts can be used to cancel outgoing amounts.

A long term application of such a single account could be to use it to reduce the burden of handling payroll tax and other administrative chores on business – the account would become the place to manage credits and deductions. Salary could be paid gross, allowances managed through the account, tax credits applied and so on. Reconciliation at the end of the tax year would be a relatively simple task because all of the flows would have passed across a single account. Such an account could be real (i.e. reflect actual debits and credits to a real bank account) or entirely virtual (an aggregation of different views held by different departments and presented as a single view to the customer). Envisioning a similar approach to businesses who need to grapple with PAYE, VAT, Corporation Tax and so on would be a no more difficult step.

The single account could be accompanied by a single department of paying money and perhaps another of receiving money. In fact though, the single account is probably not necessary to rationalise these processes – the exposure of the services from various departments through the EntArch would allow these processes to be harmonised across departments, then brought together and then fully standardised over a period.

The Citizen Data-Locker or Vault

CIP is, at first glance, principally a programme to benefit government – merging individual departmental reference numbers would provide a means to better identify people and cross-reference their activities. It could also, though, be used to support the delivery of the single account noted above. Further, it could be extended (or a separate programme could be established) along the lines of the “circumstances engine” so that it acts as a citizen-facing repository for “their data”. Probably the single biggest problem with data today is that it belongs to government, that is, it’s not in the interests of the citizen to make sure it’s up to date unless it needs to be so that they receive benefits. Moving the data that government has into the “cloud”, i.e. to the Internet arena where it can be viewed, accessed and managed by the citizen, could turn that idea on its head. Citizens charged with making sure that their data is up to date so that they receive benefits will, as a side result, make sure that other data is up to date.

The potential for reduced wastage in government here is enormous – no more sending out letters to people who have moved, no more paying benefits to accounts that are defunct, no more paying tax credits to 81 year old people without children (Sunday Mail, 24th August 2003) – note here that one of the major and otherwise hidden benefits of an EntArch is that it allows intermediaries to access government in a consistent and easily-learned way, increasing their ability to help those who might otherwise be disenfranchised.

Such a data locker has already been envisioned by the Irish government although visible progress on their programme is slight. The locker need also not be a massive Big Brother, centralist, initiative. It could easily be developed through an

opt-in process where people were presented with the chance to update data as they moved through various online services, with the results being stored (with their consent) in the vault, and therefore accessible to other parts of government – permissions could even be set (share without asking, ask me before sharing, never share and so on).

The Department of You

Fusing the rules engines, the data vault and the account those thinking far ahead may see a “Department of You” – a virtual government entity that knows who you are, what your characteristics are and one that can react to changes in your circumstances in a proactive way. The “DoY” would hold your profile and present information to you that was directly relevant based on that profile or based on similar profiles. This is, of course, the personalisation that has so long been talked about in the web world for some time. Few government sites are able to practice it because they have insufficient information about an individual or too few visitors to be able to build cross-referenced profiles. Personalisation promises much but, for government, has delivered little so far. That could change inside an EntArch.

For such a virtual department to flourish, more work would need to be done in the areas of standardising sharing of content (also a cherished goal of government, known through its technical term of “syndication”) – but this is not particularly different from many of the other areas where standardisation is required.

The “What if” Analyser

A big enough database of personal circumstances (the vault) coupled with a smart enough rules engine would allow policy makers to perform what if analysis unlike any that can be carried out today. The introduction of a new tax credit could be precisely modelled on the existing population. The effects of an aging population could be modelled and understood rapidly.