

IT Strategy

Introduction

This document gives a high-level framework which can be used in the University to guide decision-making over the next three to five years. This is the same time frame as the University and Information Services' strategic plans. The detail of what standards, technology or even approaches to particular issues such as authentication, authorisation and storage are not covered. This is intentional, as it is impractical to try and cover all the detail – since technical solutions are changing faster than the timeframe for the overall strategy. It is more effective to harness specialist expertise in specific issues as it is needed; such work should form specific policies that underpin and mesh with the overall strategy presented here.

Aim

To deliver services that support the University's strategic goals of excellence in learning and teaching, research, and commercialisation & knowledge exchange whilst exceeding user expectations

The aim puts the user experience rather than the technology at the heart of the strategy. In order to be able to exceed user expectations there has to be an element of setting or managing expectations as well as delivering great services. This is also helpful in ensuring that users do understand what those services are.

Service Characteristics

Anytime, anywhere

Services that are not constrained by time or location

*Our services need to run 24*7 so that they are available no matter what the user's working pattern or time zone. Similarly a user should be able to access the services from any location: e.g. students getting their results from an internet café or a principal investigator making spending decisions on her grant while visiting another institution. Many of the client devices will not be owned by the University and will encompass mobile devices such as phones as well as more traditional computers. Effectively this means that services will be web-delivered and require no client-side set-up, for example using virtual private networks with no network address-based restrictions. This is not possible for all services at present but is a growing need, and over time it is expected that it will become the norm.*

Easy to Use

Keep it simple and tell people about it

Making services easy to use is more important than making them functionally rich. If people cannot use a service they will invent duplicate solutions, requiring high levels of support and training and resulting in higher costs and decreased user satisfaction.

Integrated

People expect organisations and their IT to be joined up

We all use online services where we get immediate results, for example buying an airline ticket: we do not expect to be referred to different departments or to come back later to check that there really is a seat available. Traditionally the IT functions in most organisations have been delivered by monolithic applications – for example, Virtual Learning Environment, Finance, HR, Student Records and Research Archives. We now need joined up functions that cross these silos, so that – for example – a student can change course online, have the fees adjusted, make payments and get the right materials from the learning environment – all in 'real time'.

Secure	Not a concern for the user
<i>All users should be confident that information they have entered into systems will be maintained securely, will not be improperly accessed, will be secure in transmission, and will be made available to them when they want it and how they expect it to be.</i>	

Principles

Many of the principles that follow are inter-dependent, or different facets of the same issue – for example, standardisation is a mechanism for maintaining simplicity.

Keep it simple

Most services have large elements that go unused because people simply don't find them, or users are unable to understand how to use them. This has multiple implications for service providers: additional cost for no return, users inventing duplicate solutions, cost of training and support high and user satisfaction low. An important dimension that is often overlooked in the University culture is that the desire to capture all possible requirements and deliver 'perfect' solutions leads to slow solutions; a more agile approach of making a service available quickly is often a better solution than having a more complex solution delivered later.

Once a core functionality level has been met, ease of use rather than functionality is the deliverable that causes users' expectations to be exceeded. There is much evidence to show that 'core functionality' is a much lower barrier than might be expected.

User focused

Clear communication and simple routes for getting help and support are essential if we are to satisfy user expectations. There are many examples of services in the University being under-used because users, and in some cases the IT support staff, do not understand them, because the information needed to use them is not clear or easily available. Obviously where services are complex this exacerbates the problem.

De-Duplication

Duplication of services, where different parts of the organisation look to achieve similar outputs through different means, is an expensive problem that costs both to implement and to maintain; as such it is to be avoided. Often duplication is only considered in terms of duplicating centrally provided services; however, it may also occur where several Schools duplicate services that are not provided centrally. Where duplication occurs, it is usual to find that a small element of the 'duplicated' services are not common, or not delivered as effectively, and this is often the justification for the duplication. As central services are developed they will increase in functionality and many of the reasons for local duplication will become irrelevant. The problem of removing duplication through the development of appropriate central services is a governance issue that is greatly eased if the quality of the central service exceeds user expectations.

Standards

Using standards – whether they are standards we have created ourselves, industry standards or internationally accepted standards – is an important strategy for reducing complexity, removing duplication, fostering collaboration and managing relationships with vendors. The more 'standard' the item is, the greater the pressure there should be to use standard offerings. What standards are adopted and how their use is encouraged is a significant governance issue, as everyone will support the use

of standards until it comes to the crunch and they have to compromise to comply. An important element of the use of standards is having relevant policies that make relevant information easily accessible.

Technology Change

The IT industry is young, and the pace of change is high and will continue to be so for the foreseeable future. The University must retain the flexibility needed to be able to take advantage of new technology, whether that be improvements in hardware or software, new applications, new models of delivery such as open source development, software as services, outsourcing, and the general consumerisation of IT where we can all get free services from the ‘cloud’ without any IT department involvement.

Compliance

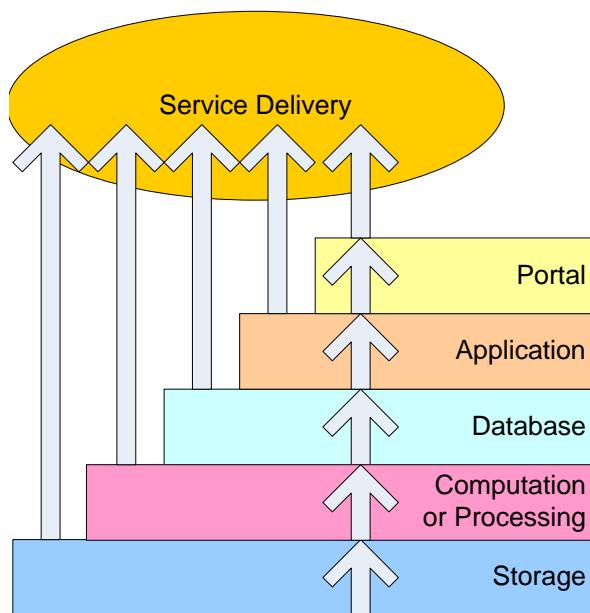
Legislative compliance, whether it is Disabilities, FOI or data protection, is an integral part of the environment.

Practical Considerations

The aim and principles lead into a number of practical considerations that are helpful to use to guide thinking towards implementation.

Layered Model

The principles lead us to the idea that increasingly complex services can be built from much simpler components that form layers. Taking as an example the delivery of administrative services, the following simple diagram illustrates the principle of the layered model:



To deliver most services we need to store data, for example for the student information service (EUCLID), University Web Site (UWS), the Library Catalogue, the Edinburgh Research Archive, etc. By standardising on the way in which we store data, a common approach to all data storage can be adopted. This has the potential advantages of reducing support costs, economies of scale in the procurement, and ensuring a common backup with similar cost reductions possible there. We are doing this with the storage area network (SAN). This is represented as the storage layer in the diagram. Within a layer there may be more than one service offering, and it is usual in a large and complex organisation for this to be the case.

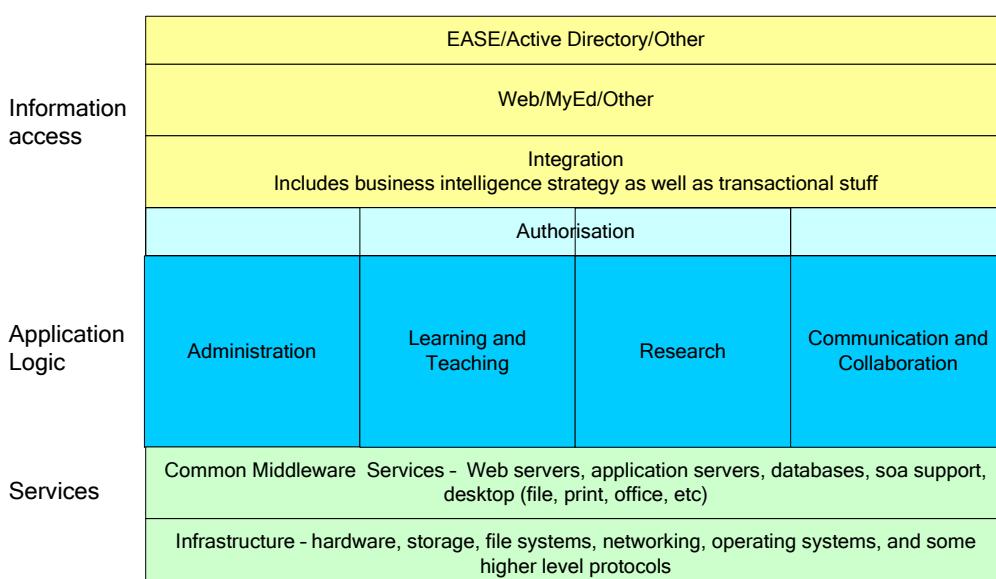
The identification of layers and the services in each layer helps to reveal unnecessary duplication.

An important aspect of the model is that a service may be exposed to end users or may be combined with other layered services to deliver higher value services. To continue the example, all of MyEd, UWS and EUCLID require data processing or computation and an underlying database. By standardising on the Oracle RDBMS running on Unix computers we are able to reduce support costs, get economies of scale – in this case by negotiating a site licence for the software – and again enabling common support services such as disaster recovery, security patching *etc*. The Oracle database service can then be exposed to applications that are written and managed by Applications Division or other staff at the University, or through open source initiatives or commercial software vendors. In an ideal world there would be only one database in the layer; however, when working with a wide range of vendors it is not practicable to have a single database platform, so the strategy adopted for centrally-delivered applications has been to use Oracle as the first choice, and MicroSoft SQL server and the open source MySQL as necessary.

In the research domain, direct access is already offered to the lower layer storage and compute-cluster services (*via* ECDF) where nearly 200 TBytes are stored already and the 12 TFlop cluster is used to capacity. At the next layer up, a University-wide global file system could be provided (*e.g.* AFS) to allow storage to be accessible from anywhere on and off campus. Another example could be a code-versioning framework which would sit above the storage layer. At higher levels both Web server and database services are already used.

Similar principles apply to learning and teaching; for example, the use of video in learning materials delivered by the VLE requires storage of the media, streaming services and presentation via applications and web interfaces. Adopting a layered model enables the same storage services to be used as in the admin example and the streaming video service to be shared with the University web presence.

A more complete but still simplified model which does not show exposure of individual services (the steps in the diagram above) is shown below:



Standardisation

Standardisation is important, as it reduces complexity – which lowers support costs and increases the scale of the service, which in turn reduces unit costs. These two drivers, reduced complexity and scale, will enable improved service and reduced costs. A great example of this is the introduction of the managed desktop for Windows PCs. However, we cannot ignore the requirements of the business in this process of standardisation. We have to be able to balance the desire to standardise to improve service and reduce cost against the need to enable our users to use the tools that are appropriate for their work. To this end we have adopted policies that define what centrally-supported services are across a limited range of choices. For example we have Windows, Macintosh and Linux for desktop PCs, Mozilla Firefox and Internet Explorer for web browsers, and Oracle, SQL Server and MySQL for databases. There will always be some people who need to operate outside the prescribed environment. It is important that we balance between appropriate pressure to conform to standards and the need to innovate. We have a few areas where control is essential, such as procurement where we have both benefits to achieve and legal obligations to comply with. In the main, the more ‘standard’ an item is, the more pressure there should be to use it; good examples would be buying a desktop PC or a standard software applications. To achieve the desired approach, we need to encourage people to use the standard route, by making it clear what is expected and by delivering great services. We should not necessarily prevent people doing what they want (it’s too hard), but we should not support them when they move beyond the central set. There is a difficulty in balancing the desire to allow people to do their own thing, with the possibility that they will adversely affect other members of the community.

No central service can keep pace with all of the new innovation. It is therefore important that the work of those at ‘the sharp end’ can, where appropriate, be brought into the central service so that the benefit of the innovation can be enjoyed by all. Good examples of this issue lie in the virtual learning and research environments where ideas, tools and services are still evolving and what constitutes a mature service has yet to be defined. At the start of the academic year, the centrally-run service based on WebCT introduced an e-Portfolio tool which is expected to become the standard offering. With this improvement to the service, all learners can have it, they only need to engage with one system, and institutional support costs are reduced. However, those areas that were innovative and introduced e-Portfolios ahead of the central service provision will have to go through some transition over time to be able to use the centrally-provided service. Processes for managing the cycle from innovation to standard service delivery are not easy, as they run into all the pressures of people and the ways they behave, together with the ability of the service provider to deliver and the need for a governance model that can differentiate between innovation and duplication.

As we engage with more and more suppliers, a problem arises because of duplication in the functionality that is provided; for example, e-Financials, e-Vision (EUCLID) and WebCT all have elements of a portal and offer elements of single sign-on that duplicate some of the functionality – but not all – in MyEd/EASE. Many vendor strategies are based on what could be described as ‘King of the Hill’: if you use my portal, my authentication, my ..., everything will work terrifically well (and it helps to shut out other vendors). Increasingly, the effort of the central service supplier has to be on integrating these services so that for example an announcement made in WebCT is delivered via MyEd and/or WebCT. This is often difficult because the vendors don’t provide the interfaces needed to allow this to happen, and the ‘King of the Hill’ strategy fits well with implementation projects where there is a single focus. Again this produces a governance tension.

Service Orientated Architecture – SOA

The SOA approach fits well with the strategy, as it is effectively a mechanism to deliver business logic that follows the principles of the layered service model. SOA is very much in vogue at the time of writing: it has been adopted by the major vendors of business applications SAP, Oracle, Microsoft, IBM, etc., and it will play an important part in both industry and University strategies. This section is included because the way in which SOA works illustrates many of the principles that we wish to adopt in the University, and it will form an important implementation strand over the next five years.

As the SOA approach may not be familiar to all readers the following is a detailed but still simple explanation.

Service orientated architecture – SOA – is a way of disaggregating the functionality in a large business application into what are called *services*. ‘Service’ is a much used word in the IT industry and means many things to many people. In this context, services are business functions which are generally much smaller than the functionality in a whole application. A couple of examples follow giving both an education perspective and a large-scale commercial application – demonstrating both its applicability to HE and its ability to deliver in very large scale applications

Example 1 – there are an increasing number of tools in the collaboration space: blogs, wikis, discussion fora, email, ediary, etc. In many cases, users want to be able to work with a set of collaborators, whether it is their research colleagues who may be spread around the world or groups of students they are teaching. Adopting a central group management service, as has been done at Newcastle University, and exposing its ability to update groups and to publish groups as services, allows a mechanism for applications to use those services to enable real time synchronisation of groups across multiple tools. The user only needs to update the group once but can immediately use all the tools with the updated group. The advantage from an IT perspective is that the services only have to be published once and then many applications can use them, rather than having to build point-to-point integration for each tool combination.

Example 2 – the DVLA application that enables the public to buy road tax online needs to check whether the vehicle has an MoT and insurance. These checks are done using services that allow the DVLA to send a request to all of the participating insurance company systems and their MoT system and get a response in ‘real time’ that enables the end user to complete the online application and payment. It does not matter what application the insurance company is using, so long as it can deliver the right service interface to the DVLA. The services provide a mechanism to deliver a great customer experience, as the process seamlessly integrates the MoT and insurance check. From the DVLA point of view, it only needs one service for all insurance companies – not one per company as would have been the case a few years ago, greatly reducing duplication and hence its costs.

The services are based on standards that are independent of the hardware/software platform being used, and are loosely coupled together via the Internet. The use of standards gives us the flexibility to extend the functionality of applications with components from more than one vendor or in-house developments, the ability to provide cross-application integration in real time, and an easier mechanism to replace components with new services as they become available.

The SOA approach fits well with the layered model as it is effectively disaggregating the business logic layer into smaller components or layers. It enables us to avoid duplication, as we can provide a service like ‘send announcement’ that may be called from many locations, for example in the portal, in WebCT, in a workflow, *etc*. This addresses the ‘King of the Hill’ problem, allowing us greater control over the user experience rather than being locked in to an individual vendor’s approach. Similarly it provides a framework for Schools to extend centrally-provided services to meet their local requirements. The ability to extend functionality in this way using a framework based on standards will encourage innovation and agility whilst retaining the necessary control to deliver robust services.

As SOA follows the layered model, the governance issues are very similar to those of the overall IT strategy: who is allowed to use services, how do you avoid duplication, overlapping but slightly dissimilar functionality, dependence on others for availability and difficulty of problem identification/location.

Transparency and Charging

As the technology changes there is a constant need to review allocation of funding so that new services can be implemented. However, funding based on historical allocations restricts the flexibility to change. It is extremely difficult to withdraw existing services, so change is limited to the amount that can be committed from the central allocation to new initiatives and the amount that can be saved from running costs of existing services – working smarter, cheaper technology *etc*. Generally, the amount of money available to invest in new services will be limited and will always be insufficient to fund major new initiatives such as EUCLID or the University Web Site. In these cases funding has to be requested from central funds if the projects are to proceed. It is just as important to stop doing things that are no longer required as it is to ensure that new initiatives are not allowed to proceed if funding restrictions mean that user expectations cannot be met.

In some cases the University has taken the view that funding for particular initiatives should be provided through charging at the point of delivery, so that the service can be scaled to meet demand: for example, charging for network ports and SAN disk usage. Charges of this kind can be very helpful, both in making sure that new services are adequately funded to deliver the required quality, and in generating a more entrepreneurial and agile environment whilst stimulating greater accountability.

Given that we are unlikely to have a radical shift in our funding model, it would seem sensible to adopt a mixed model which aims to deliver some services to users where all costs are met from central provision (for example, Finance, HR, Library Catalogue), some where a mixed model applies (for example Network, where there are some charges aimed at recovering the marginal costs of installing and running new ports and wireless access points), and some where there is full cost recovery (for example delivering services to PIs, as they are expected to recover full economic costs from their funders or work for external bodies). In all cases, whether the service is being delivered with or without charges, we need to be able to determine costs of service delivery transparently.

The mix of funding types, charges, withdrawal of services, *etc* are areas where central service directors would be expected to seek guidance from the governance bodies supporting IT at the University.

Partnership

As described earlier, no one area of the University can do it all – whether that is central IT provision from Information Services or research group based computing. We all have symbiotic relationships and many partnerships to sustain if we are going to achieve the best. This may be internally, within the University, and/or with partners and collaborators in other universities or research institutes, or indeed with commercial vendors such as Cray with Hector or Tribal with EUCLID. Respecting, encouraging and developing these partnerships are important to our combined success, and form an important part of the strategy. Providing our approach to standards and duplication is followed, there is no definition in the strategy that defines who should fulfil a particular service. (Peter very happy for more in here)

Governance

There is a need for a place where strategy can be developed and the kinds of governance decisions about standardisation, duplication, authorisation, charging, *etc.* can be brought. The first Kenway Review recommended that a Group should be set up to do this. In many ways this would be a re-formulation of the IT Committee. The Group would be advisory to the Vice Principal (CIO role) and formally report to the Knowledge Strategy Committee (KSC). The overall strategy would be signed off by the KSC, but more operational policy would be developed by the Group and short life working parties as necessary.