Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

Transformation Management Services (TMS) presented a response to the Productivity Commission's draft report in Melbourne on the 10th of June 2014.

This presentation included a proposal in response to the questions the Commission had raised in relation to the information technology (IT) costs of Australia's justice systems. The proposed concept was that the Commonwealth sponsor a project that would address the Australia–wide need to reduce Justice IT costs and provide a coherent framework for monitoring and reporting justice outcomes and for modern case management.

Commission Question:

The Commissioner asked TMS (Michael Hall) to respond with an estimate of the costs of mounting such a project and queried the Attorney-General's capacity to auspice it.

Response from Transformation Management Services: Michael Hall.

The TMS project cost estimate for the Commission, on the day, was about $20 to $25 million. I confirm that cost in this more detailed estimate and provide the explanation for my thinking below.

Understanding of the estimate requires an explanation of the nature of the proposal, the scope of the work and appreciation of potential benefits in relation to alternative systems and their costs. Unfortunately some parts of this necessarily reference uncommon industry products or terms. I have referenced accessible online resources for these where appropriate.

This estimate is brief and does not go into detail on many of the aspects of the system design I referred to in my oral submission to the Commission.

Background - Necessary Understanding

Current model of funding IT Justice is too expensive costing $140m Australia-wide every 4 years

Modern approaches to justice and modern case management expectations are hampered by inflexible IT systems built to respond to justice needs that are decades old. The approaches taken to develop these systems do not now represent best practice in scalable architecture, database design, development environments or integration with reporting tools that analyse data trends.

Government justice departments in all jurisdictions are wary of undertaking development. Instead they look to purchase off-the-shelf systems and assume that large vendors, typically with global IT businesses, will provide protection from risk. In this assumption they have not been proved correct (cf. Victoria - Justice IT Project).

Our estimate of the cost of redevelopment of a major court/tribunal system using the technologies popular for the last 15 years is around 7 to 15 million dollars each. [i] [ii]   A quick scan of the main justice systems in each Australian jurisdiction shows there to be approximately 8 tiers of systems in 5 major states (40 systems)and 6 tiers of systems in 3 minor (by population) jurisdictions (Attachment 1).

With a typical life of about 12 years, it is not an heroic assumption to place one third of these systems as requiring redevelopment or major revision at any one time. Over a the next 5 years we could expect Australian governments to spend about $140m every 4 years on Justice IT systems development.

While some of these systems (Vic Supreme Court) have been replaced recently, the result of an initial $40m court IT funding allocation, others are in the process of review (SA , NSW) or were to

have been replaced but are waiting commitment of funds or project approval or are in the process of redevelopment (Vic).

Transformation Management Services suggest a different approach to that of accepting the renewal and redevelopment of each individual system for each jurisdiction.

Rationale and Proposed Approach

Key value database systems can address many of the shortcomings of traditional relational database systems and for the first time provide an opportunity to develop an Australia wide Justice IT system that can be 'grown organically', to meet all major case management, reporting and client application expectations.

Transformation Systems (a subsidiary of TMS) has extensive experience in developing large scale transaction management systems[iii]. Many of these would be described as case management systems with integrated reporting, auditing, alerts and exception monitoring.   All of the earliest of these were based on relational data models and architecture. Later models adapted significant de-normalisation of data to provide access speed and eventually evolved into a key-value architecture, relational hybrid model.[iv]

Database architectures are a large and sometimes controversial topic. Our experience of this and knowledge of what is happening in the wider online, virtual server and cloud environments lead us to believe that this is the architecture of the future for transaction and case management systems. We believe that the time is appropriate to assess this environment for its ability to provide solutions to IT justice management issues that existing case management, single court systems have not been able to do.

The three components I think are most important to the development of new justice systems are:

- **Speed and responsiveness of the key value architecture for large scale data manipulation**

- **Scalability of a distributed – clustered environment for wide connectivity**

- **Engagement of a motivated developer community through a published open source API**

The development of large scale databases and the applications that access them has changed in the last 4 years. The major developers of online systems (Google, Amazon, Facebook, Netflix etc and many fortune 500 companies) are employing 'Key-value' flat-table, distributed data models to manage access to millions of records and web pages in real time with sub-second responses.

Traditional relational database models cannot match this performance. As their record size and data complexity grows, the page locking required to add new records and rebuild indexes becomes debilitating. Even employing bigger, faster computers with more memory resources cannot provide the performance users require from old data models.

With the success of the large online organisations and their key-value database systems there has been a growth in business interest in large scale key-value systems derived from the same architecture. Google has now made most of its intellectual effort in developing 'big tables' and clustered scalable computing available for scrutiny. Many large companies like Google have developed Application Program Interfaces (API) and provided these through open source licensing that now gives developers around the world free access to their architecture and the opportunity to build commercial applications that interface with their systems.

Projects in collaboration with open source communities and Universities have delivered what are now business usable large scale data management architectures. Facebook provided the nucleus for

the Cassandra project[v]. This is the most popular of the key-value clustered data bases. There are many others.

The open source community is a major engine of world IT development. Ohloh[vi] reports on over 660,000 current projects across all software communities amounting to over 21 billion lines of open source code.

It is the new development communities, the open source model and the possibility that this mechanism can provide a commercial basis for development of most of the front end applications (Apps) that an integrated Justice IT system will require that makes the third component of development attractive.

Moving to new data bases requires new data models

There are fundamental differences between the data modes used in traditional relational databases and key-value data models in terms of how they organise data and use it. With the flat databases there is much more focus on 'Mapping' (filtering and sorting records) to support the potential reporting requirements and 'Reducing' collecting aggregate measures (counts and totals) in advance of their use. Hadoop is one of these 'big data' management engines.[vii] This MapReduce model helps process massive amounts of data in parallel and it is becoming pervasive. It allows data from large clusters of computers to be managed and quickly deployed in user queries[viii].

Full discussion of the approach would be too lengthy here. It is perhaps sufficient to say that a fundamental rethink of the collection, storage, indexing and aggregating of data is required to apply new models. The data objects themselves: names, documents, cases, transactions, events do not change but their inter-relations with other data and the way they are stored and accessed and the computer servers they use do change.

Conclusion

The new architectures and open source approach can be applied here, now

My conclusion is that the new data approach developed by large online organisations is sufficiently mature and accessible and the open source model sufficiently trialled and understood to be used by government in a large scale justice IT systems development.

What a national justice IT system would look like and how it would work

Local justice systems would adopt common basic applications developed nationally and then adapt them, using local IT programmers. Commercial developers would develop applications that could interact with the national system using the common elements and published API; and with police or legal practitioner IT systems and which could be sold locally.   The costs would reduce with commercial activity 'unbundled' from large single provider systems to a range of smaller tailored systems and even to some simple mobile apps. Competition would generate savings.

The major benefit would be a single project definition of data objects. This would provide a basis for case management, reporting and resource management. It would also integrate with private resources building specific purpose applications required by justice workers and the public and peculiar to local justice practices.

On the hardware side, the new data structures offer the opportunity of highly scalable, distributed data clusters that will self-repair and are extensible from government to large and small organisations.

The next sections provide a brief consideration of the scope of the project required to implement this proposal and the phases of work involved.

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

Project to Develop Australia's Justice System Architecture

The proposal is to establish the data model for justice systems first as information definitions and exchange setting out a common basis for Justice transactions and the reporting capacity required by jurisdictions. Secondly it would establish an online environment where client systems will be developed that will integrate with this in a commercially viable model.

There are five components to the estimated costing ($m)

1.	A scoping phase	2.5
2.	A data gathering phase	3.0
3.	A structural development phase	7.0
4.	An API infrastructure Phase	4.0
5.	A pilot app development phase	6.0
	▪ Grants	3.0
Total estimate		**25.5**

1 Scoping

The initial task is to determine the viability of the model. To a certain extent this is a nation building project and will depend on the persuasion of the auspice agency. The initial project must crunch the numbers on benefits and costs and lay these before the jurisdictions to obtain majority support.

This is not so much a government to government agreement but a set of Court to Executive government agreements. While governments may require dispute management arms and tribunal to cooperate, success will require the active co-operation of the Courts. Visits to major jurisdictions will be required. A report on the current state of Court system development and renewal will be part of the analysis.

Scoping will require several experts in the architecture sourced from industry, operational courts management expertise and staff familiar with separation of powers issues and executive government. It should be located in the Federal Attorney General's department.

I estimate the staff required as 10 and the time required at 12 months. The cost is estimated at $2.5m   (salaries $1.4 on-costs $600k, recurrent $500k)

Scoping will necessarily require the development of an independent portal to manage interactions with potential participants and include gathering base information on contacts and current systems.

This could also include the testing of various data models, classes of data mappings and especially the identification of any 'show stopping' problems including jurisdictional or legal problems.

2   Data definitions and modelling transactions and the data exchange domain

On successful scoping – that is a determination to proceed with a full development of the proposed justice IT system, detailed gathering of data will be required. This includes:

- The development of an online portal for managing data gathering and verification.

- Identification of definitional data for future database columns or keys,

- the size dimensions of expected data sets, document and repository loads

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

- The relationship between providers of data and their access to the system to input data

- The requirements for reports (a generic term for all system outputs)

- Initial security considerations: encryption and access.

- The consequential hardware requirements for initial system piloting.

Existing court systems in all jurisdictions will provide most of the input for data keys (columns) and their definitions. A good part of the initial work will be liaison with the jurisdictions gathering their definitional data.

It should be noted that this phase is not about getting agreement on definitions. The task is to collect and align definitions in an underlying key or column or object-based architecture. For example a name is always a name and a transaction is always a transaction. The position of a family name on an application page (ie for Chinese names) is a local display issue not a database definition issue. This means that definitions may be employed in the various applications in each jurisdiction to meet local expectations. The architecture should also allow multi-language presentation an aspect which has been difficult in traditional systems approaches.

The result of the data gathering will be a project outline and a light treatment of the functional specification, bearing in mind that the eventual development will be open source and limiting the functional expectations in the initial phases of a project is antithetical to this development environment. In essence this is the phase where all justice data is a potential input for the system. Later more ruthless decisions will be made as to which data sets are important and which can be relegated to non-core activities.

The functional outline of core case management components of the system will need to be endorsed by the various courts and by governments on behalf of the lower tribunals and dispute bodies.

The outputs are online schema, available for interrogation by courts, tribunals and justice stakeholders.

Estimated staff required 12. Time required 8 months. The estimated cost is $3m on a 12 months project basis, although some functions may overlap with initial scoping and later development phases.

3   Structural Development and data models

Structural development includes the initial software engineering phase for the development of the software environment, the data repository, software classes and code libraries defining the API requirements in broad outline and in its expected client applications and interfaces.

The open source industry has developed standardised componentry for managing data base interactions. An Initial component of this phase is the establishment of the software stack that will become the basis for the software environment.

The demands of the database environment will determine the hardware required. Our expectation is that this will follow a model of clustered peer servers of much lower cost than mainframes used in the past and providing low cost scalability and redundancy. This phase will require the establishment of cluster of computers to model the justice IT system and trial interactions with other clusters and clients. Input from the business community and projects using leading key value databases like Cassandra will inform the decisions on the database environment.

The major work is in identifying and defining justice transactions in terms of the data model objects, the classes providing the functionality for interactions with the object model and the demands that data aggregation and reporting will make. To a large extent the development of the models are iterative, depending on a trial and testing then integration with other components and up-scaling.

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

Our expectation is that with the right software engineers exposed to leading edge key value database projects the report models will be designed in a modern parallel processing model (see MapReduce above). Output from this phase is a well-defined core data structure, class definitions and code libraries for interaction with the database repositories. The database environment itself will provide the scalability and cluster management for integration of private and commercially managed clusters.

Input from key government and industry stakeholders would also commence in this phase. Minor projects may be agreed that support commercial development of client applications that will interact with the Justice IT system. These include case preparation, case risk assessment, document management and scheduling applications.

The project management would be continued from Phase 2 and to an extent may overlap in time with that phase and the later API development.

Project management and roles are beyond the scope of this brief estimate.   With justice IT being a prospective project, care will be needed to define client/user expectations to stand in place of existing clients. The suggest development approach is agile and iterative.

Estimated costs for this phase are around $6.8m. The period of time required will be 12 to 18 months. Project executive salary costs are part shared with other overlapping phases.

4   API infrastructure Creation and Testing

The development of a successful API will require engagement of client groups, specific agencies and businesses willing to commit time and funds to the development of applications that interface with a Justice System. This is the creation of a software environment that is both technically sufficient for completing transactions, meets security requirements and developer expectations for extensibility in the future. Towards the end of the project, components of the API may well be handed over to the developer community of this is sufficiently mature.

The API is: a portal, a project management website and a ticketing system for jobs. Sufficient software engineers will be required to rapidly develop data classes, write code, modify the specified database structures and bring these together in an attractive package that will encourage commercial extension of the core justice functions.

The bulk of this work is software engineering. I estimate 5 engineers would be required with a range of skills and interests plus technical support and testing personnel. The mobility of these groups and the industry demand will make attraction and retention problematic. I have estimated an uplift of 30% on top of usual costs to recruit and staff this phase as these skills are sought after.

The output of this phase is an environment that supports the development of applications which link to the data repository and potentially other private justice data resources. The environment must include the capacity to be dynamically improved, to retain backwards compatibility in fact all aspects of a mature open source API.

Estimated costs are $4.3m (a 12 month cost) although the period to initial public exposure is estimated at 9 months with further development taking up to another 6 months.

5   Development of pilot applications and promotion of data security

Pilot applications can be developed in house and by the Client stakeholders. The essential aspect of an open source development environment is to encourage use of common components. It permits their development as long as the improvements are placed fully in the open source environment for further use. Against this we have a need to encourage private investment in the Justice IT system and must acknowledge the stakeholder's right to differentiate their products and services for their clients.

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

Pilot applications funded or part funded by government should remain open for use by anyone whether developed by private stakeholders or not. Where stakeholders use core code or components in a unique application they are entitled to retain this for their own use or to sell licenses to the application they have developed – having regard to the fact that it is only their code they are able to license.

Application development in-house would include case a management for at least 2 court models – civil and criminal and one conciliation mediation model (non-adversarial or inquisitorial).

These can be expected to cost in the region of $2m each given the data base components are already functionally specified by earlier phases and the API has been defined. A total of $6m is estimated for in-house development of the front-end aspects of 3 case management, reporting and analysis systems.

A further estimated $3m including admin should be allocated as grants to partially support development of private client applications on the basis that the code base of these applications will be open source and free for use by other clients.

Total estimate for the pilot application development phase is therefore $9m.

The total estimated cost of the Justice project is: $25.5m[ix]

This would provide a starting infrastructure for the data repository for federal case management and case management systems for 2 justice streams and one inquisitorial tribunal model. These would also be applicable to state case management with some minor modifications.

Expansion of the system to provide hardware clusters to support state jurisdictions would require additional costs.

Michael Hall
Principal
Transformation Management Services 2014

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

Comments and additional considerations

Our experience with justice case management systems indicates that most of the data objects may be captured in similar and generic ways amenable to a key value model.

1   Security concerns and layering security

Concern for encryption and security of data and access are not an overlay to the justice system but an integral part of design and development. They are a day to day concern for analysts and developers and their consideration is included in the development costs.

2 Application Licensing

The developed stakeholder applications will require additional legal scrutiny of licenses. This may be provided by the AG's department or contracted at additional cost. For example, nothing should prevent other stakeholders developing similar or functionally identical applications that exploit the central justice data repository in the same way – ie no restriction on look can feel for client applications.

3   Location of justice system project

The commissioner questioned whether the Federal Attorney General's department was the appropriate auspice agency to manage the project.

The AGs is already the lead agency for government's efforts in tracking and combatting Computer Crime (CERT)

While the broad range of justice responsibilities is greater in the states it would seem that a Federal Management of the project would be required to gain the trust and co-ordination needed for its successful completion. Given that it is likely that at least one of the major jurisdictions and a University may well be interested in providing resources. I would also encourage the wider liaison with major law firms and industry representative bodies to represent the interests of the bulk of solicitors and legal workers.

4   Location of Hardware (servers)

The concept underpinning the justice IT system is a highly extensible redundant architecture. The core concepts are clusters of self-managing or self-replicating data repositories (nodes) that balance loads depending on demand and will interact with new clusters managed in each jurisdiction as they are brought on-line to support additional client work loads.

This type of architecture is well suited for the virtual computing environment. The costs of managing the computing resources in government data centres vs commercial data centres would be expected to determine their eventual location.

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

State and Federal Justice System estimated renewal costs   **Attachment 1**

States Justice systems:

Complaint, Ombudsman, Mediation, Conciliation, Tribunal, Local / Magistrate, District, Supreme.

Major population jurisdictions: NSW, Vic, Qld WA, SA. Minor: Tas, ACT, NT

Totals: major- 40(8 x 5) and minor - 18(6 x 3) = 58.

Federal Justice systems:

Complaints, Ombudsman, Conciliation(various) , AAT, Fed Magistrates, Federal Court, Full Bench

Total major - 7

Also Included are: AATs, Children's Courts, Specialist Courts, Court Annexed mediation

Progressive redevelopment costs for Australian Justice Systems:

Expected 12 year system life. One third under major revision or redevelopment at any one time.

Total capital development: 47 x$7m ($329m) + 18 x $5m ($90m) = $419m

One third redeveloped every 4 years = $136m.

Notes:

Capital/development costs - Operating costs are excluded.

Not within scope: Courts of Appeal, High Court, Family Court.

Michael Hall – Transformation Management Services   2014

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

Example project component and phase resources and estimates. **Attachment 2**

(First 3 phases attached here - All phase costs are in an online document accessed from the end notes**)**

1 Scoping Study				($000)	
1	Project Director	250	250		
1	Deputy Court	160	160		
1	Industry specialist	220	220		
2	Software Specialists	150	300		
1	Court Liaison	100	100		
1	Court research	100	100		
2	Researcher	90	180		
1	Project Manager	120	120		
1	Support	80	80		
	sub total			1,510	
	on costs @ 40%		604		
	Fixed Costs @25%		378		
	sub total			982	
	Total Scoping Estimate				**2,492**

2 Data Gathering ($000) **$k**

2 Data Gathering					
1	Project Director	250	250		
1	Deputy	160	160		
1	Portal Software	130	130		
1	Court systems	150	150		
1	Tribunal Systems	140	140		
1	Dispute Systems	130	130		
1	Project Management	180	180		
4	Software specialist	130	520		
1	Tech Support	100	100		
1	Support	80	80		
	sub total			1,840	
	on-costs @40%		736		
	fixed costs @25%		460		
	subtotal			1,196	
	Total Data Gathering				**3,036**

Productivity Commission Access to Justice

Estimate of the cost of an open Australian Justice open data scheme

3 Structural development ($000)							
1	Director	(half)	250	125			
1	Deputy	(half)	160	80			
1	support	(half)	80	40			
					245		
Civil							
1	Analyst		180	180			
2	Lead Prog.		140	280			
4	Member		100	400			
3	Tech Support/test		80	240			
					1,100		
Criminal							
1	Analyst		180	180			
2	Lead Prog.		140	280			
3	Member		100	300			
4	Tech support/test		80	320			
					1,080		
Project Management							
1	Project Manager		180	180			
2	Coordinator		120	240			
3	Consultant		100	300			
2	Documentation		120	240			
					960		
	sub total Sal.					**3,385**	
	on-costs @40%				1,354		
	Fixed costs @25%				846		
	sub total					**2,200**	
	Total Project Recurrent						**5,585**
	Software Lic. and Lease				400		
	Hardware and coms				800		
	total capital					**1,200**	
	Total Structural Phase Cost						**6785.25**

References

i http://www.adelaidenow.com.au/news/south-australia/sa-government-says-cost-of-upgrading-justice-systems-archaic-it-network-will-exceed-50-million/story-fni6uo1m-1226796907236
ii http://www.lawlink.nsw.gov.au/lawlink/Corporate/ll_corporate.nsf/vwFiles/MR_14_81_million_technology_upgrade.pdf/$file/MR_14_81_million_technology_upgrade.pdf
iii http://www.transformation.com.au/home/home.php
iv http://en.wikipedia.org/wiki/NoSQL
v http://en.wikipedia.org/wiki/Apache_Cassandra
vi http://www.ohloh.net/
vii http://en.wikipedia.org/wiki/Apache_Hadoop
viii http://en.wikipedia.org/wiki/MapReduce

ix Detailed project phase costs are set out in the Google Docs Spreadsheet at this URL
https://docs.google.com/spreadsheets/d/1Jm6mUWMtETlUwG1EIHrilSyGm8H4f6cb00lJ8lMN0a0/edit#gid=0

Michael Hall – Transformation Management Services   2014