



JISC Final Report

Project Document Cover Sheet

Project Information			
Project Acronym	Shintau		
Project Title	Shib-Grid Integrated Authorization		
Start Date	1 March 2007	End Date	31 July 2009
Lead Institution	University of Kent		
Project Director	Professor David Chadwick		
Project Manager & contact details	Professor David Chadwick University of Kent, Computing Laboratory, Canterbury, CT2 7NF. Email: d.w.chadwick@kent.ac.uk Mobile: +44 77 96 44 7184		
Partner Institutions	Internet 2		
Project Web URL	http://sec.cs.kent.ac.uk/shintau/		
Programme Name (and number)	e-Infrastructure (security)		
Programme Manager	Christopher Brown		

Document Name			
Document Title	Final Report		
Reporting Period	March 2007- July 2009		
Author(s) & project role	D.W.Chadwick, Project Manager		
Date	14 Oct. 09	Filename	finalreport.doc
URL	http://sec.cs.kent.ac.uk/shintau/finalreport.pdf		
Access	<input type="checkbox"/> Project and JISC internal	<input checked="" type="checkbox"/> General dissemination	

Document History		
Version	Date	Comments
0.1	12 Oct. 09	Initial version
1.0	17 Oct. 09	Final version

Shintau Project Final Report

Table of Contents

Shintau Project Final Report	2
Table of Contents	2
Acknowledgements	2
Executive Summary	3
Background	4
Aims and Objectives	4
Methodology	5
Implementation	6
Outputs and Results	7
Outcomes and Impact	8
Conclusions & Recommendations	9
Implications for the future	10
References	10
Appendix 1 Requirement Questionnaire Recipients	11

Acknowledgements

The University of Kent would like to thank JISC for its funding of the Shintau project under its e-Infrastructure (Security) Programme. We would also like to thank Nate Klingenstein from Internet2 for his great help at the start of this project, as well as members of TERENA EMC2, Switch, Internet2, Red-IRIS and other unknown respondents who contributed to our user requirements survey and gave us valuable feedback on our initial design specifications. Without your input we would not have been able to produce the attribute aggregating system that we have done. We would also like to thank members of the Liberty Alliance group who provided valuable feedback on our initial protocol mappings and for suggesting the one that we have finally implemented which is based on the LA Discovery Service. We would especially like to thank Connor Cahill for providing his open source Discovery Service implementation upon which our linking service is built.

Executive Summary

Today we all possess a whole range of attributes issued by different authorities for example, our qualifications, club memberships, credit cards, employment details etc. When these attributes are used in authorising our use of electronic resources, we hit a big problem. Today's identity management and grid systems are usually only capable of allowing us to present our attributes issued by a single authority. This is insufficient. The Shintau project's main objective was to solve this problem in a privacy preserving way, by allowing the user to present his attributes from a whole set of authorities, of his own choosing, without letting any of these authorities know about the user's accounts and attributes at any of the other authorities. We call this process privacy preserving attribute aggregation.

The project has been remarkably successful. Highlights of the project's outputs include:

- We are on the cover of May 2009's edition of IEEE Computer magazine in a special edition on Identity Management. This article describes the Shintau model and implementation.
- We have a detailed paper describing the Shintau conceptual model accepted for Next Generation Computer Systems journal.
- We successfully collected user requirements from a large and distributed user group, and presented a paper describing these user requirements at the E-Portfolio conference in Maastricht in 2008
- We successfully demonstrated the Shintau software at the Internet2 Spring meeting and the Terena Networking Conference 2009 in Malaga.
- A live demonstration of the Shintau software is available for public use at <http://issrg-beta.cs.kent.ac.uk:8080/loademo.html>
- We have secured further funding (€1M) to continue our research into attribute aggregation and federated identity management as part of the EC TAS3 Integrated Project which runs to end of Dec 2011. We plan to add support for attribute aggregation to Microsoft's CardSpace in the next version of our software.

In order to have a wide impact on the user community in the medium term, it is essential that the following issues that we uncovered in this project, are dealt with in the short term:

1. The Shintau model of user interaction is similar in concept to the Shibboleth model in which the users establish their own Attribute Release Policies prior to engaging in the use of federated services. To date, few if any UK Federation sites allow their users to actively provide consent through these policies. Consequently users are still not familiar with this concept and have no experience of using Attribute Release Policies. We therefore currently have little or no knowledge about the users' perceptions of, or level of acceptance of, Attribute Release Policies.
2. In order to protect the user's privacy, in Shintau the user has to establish his own Account Release Policy at the Account Linking Service before any of his linked accounts will be released to the federated services. Account Release Policies are similar in concept to Attribute Release Policies but operates at a different level of granularity to them. But until user's become familiar with Attribute Release Policies they are likely to find the concept of Account Release Policies quite difficult to grasp. However we anticipate that once users do become familiar with Attribute Release Policies, then understanding, setting and using Account Release Policies will be a relatively easy next step for them to take.
3. The Shintau software requires minor modifications to the use of existing Shibboleth protocols, therefore until these changes are integrated into the existing Shibboleth and SAMLv2 software implementations, neither attribute authorities/identity providers or service providers will be able to use it.
4. The Shibboleth and UK Federation MetaData is an impediment to the ease of use of the Shintau software. Specifically, it needs to have a user friendly display name for service providers added to it, in the same way that it already contains a similar friendly name for identity providers.

In conclusion, the Shintau project has been very successful from an e-infrastructure security research perspective. It has been truly trailblazing in its results and outputs. Attribute aggregation is currently a

hot topic that is actively being discussed by the academic networking community, and we have been the first to produce a fully working demonstration and open source software. Of course, being at the leading edge means that the project's outputs are not yet being used by end users. However we anticipate that in another 5 years or so attribute aggregation will be widely rolled out to the end user community (considering how long it has taken for Shibboleth to be rolled out into operational services, this is not a long period of time). We believe that Shintau forms a solid rock on which to build future attribute aggregation efforts.

Background

Today's federated identity management systems use attribute based access controls. Grid systems are also migrating to this model, with an increasing use of XACML [5] and PERMIS [6] Policy Decision Points (PDPs). Integration of grids and Shibboleth, and use of Shibboleth for more security demanding applications is being hampered because a user's attributes are typically issued by multiple authorities, and are held in different locations under different identifiers (cf the situation with plastic cards today). Currently there is no coherent way of collecting all these attributes together and validating that they all belong to the same user so that they can be used for authorization of the user's request. The typical state of the art today in Federations is that a user's attributes are picked up from a single source, the Identity Provider, and passed to the Service Provider, where it may merge these with any local user attributes that it holds. The state of the art in grids is that a user's VOMS attributes [7] may be pushed with the proxy certificate, and intermediate GT4 GridShib code may pick up the user's attributes from a Shibboleth IdP [8]. However, no general purpose solution exists for arbitrary attribute aggregation from multiple attribute authorities (AAs). This is the primary objective of the Shintau project. The benefits of this will be huge. This is because a user typically has only one or two attributes from each AA/IdP (cf today's plastic cards). Thus in any Internet transaction or grid job a user may wish to combine the use of several attributes from different IdPs, and the set of attributes might be different for each transaction. Shintau aims to achieve this objective within the current Shibboleth and grid infrastructures through the process of privacy preserving attribute aggregation. In the Shintau model, the user is the only person who knows where all his attributes are, and which AAs issue them. Consequently the user is the only person capable of aggregating his attributes, and he is given full control of this functionality.

Aims and Objectives

The original aims and objectives are listed below. Any changes are made in *italics*, and the results are in **bold**.

The first objective of this project is to work with the international community, primarily the Internet2 consortium and the Globus Consortium, but including SWITCH, TERENA and others, to develop the Shibboleth protocol specifications, based on SAMLv2 and other protocols, that will allow a Shibboleth service provider (SP) to collect together a user's attributes from multiple authorities, whilst preserving the user's privacy, so that the aggregated attributes can be used to authorise the user's request. This will significantly ease the integration of Shibboleth with grids. However, the resulting attribute aggregation protocol will be of benefit to any Shibboleth enabled SP be it a web service, a grid service, or a conventional Shibboleth SP etc.

This objective has been achieved as originally planned.

The second objective is to build a Policy Information Point (the nAA-PIP) that will evaluate the collected attributes (or credentials) according to the configured trust policy of the Service Provider (SP) and will return the valid set of attributes to the SP's Policy Enforcement Point (PEP). The PEP can then pass the complete set of validated and aggregated attributes to a conventional Policy Decision Point (PDP) for it to make access control decisions. The nAA-PIP will be fully standards conformant, and will be called by the SP through either the standard web services protocol that is being defined by the OGSA-Authz WG [8] or by a Java API that is already implemented in Globus Toolkit and is also due to be published by the OGSA-Authz WG.

This objective has been fully met in the design specification. Only minor enhancements are needed to the existing SAMLv2 and Liberty Alliance protocols in order to achieve attribute aggregation

The third objective is to implement the aggregation protocol within the nAA-PIP so that it is capable of collecting the attributes itself from the multiple authorities, prior to validation.

This has been implemented in the Shintau open source code

The fourth objective is to build a pilot demonstrator for the National Grid Service that will show how attributes from multiple AAs can be integrated together and used in authorisation decision making at grid sites that use Shibboleth IdPs.

This objective has been changed to

- i) mount a live Shibboleth based demonstration service on Kent's web site and*
- ii) run user trials with a demonstration grid service at the National e-Science Centre in Glasgow*

Objective i) has been achieved and

ii) code and documentation have been provided to NeSC for them to mount a demonstration service for user trials, but unfortunately by the project completion date the trials had not been completed. Nevertheless these trials will still go ahead and the results will be documented, albeit not by the original planned completion date.

The final objective is to release all the developed software as open source code through NMI/OMII to the community at large, with a full set of specifications and documentation.

At the time of writing this is still the objective but it has not been achieved yet as we are waiting for the documentation and software to be validated by NeSC running a demonstration service and user trials.

Methodology

The initial project task was to gather user requirements from a wide set of grid and Shibboleth Internet users. This was achieved by means of a questionnaire. The draft questionnaire was initially tested using six participants. The results of the tests allowed us to hone the questionnaire, by removing any ambiguities and lack of clarity and ensuring that we got full coverage of the requirements. The questionnaire was then distributed widely via a dozen international mailing lists. We analysed the results to obtain the user requirements (12 were identified). We analysed the requirements against four possible known attribute aggregation models, to see which had the best fit. We produced our high level conceptual design using the best model that satisfied the vast majority of the user requirements (11 out of 12).

The conceptual design was circulated to the user community for comment and revision. Once we had a stable conceptual design, we then mapped it into existing standard protocols. Several mappings were possible and we documented three of them. We took these protocol mappings back to the user community for comment, allowing them to express their preferences for the various mappings. Finally we took the preferred mapping to the Liberty Alliance standard's group for comment, since each of the mappings required minor modifications to be made to the standard protocols. The Liberty Alliance group suggested that we use an alternative protocol mapping which required less enhancements to the standard protocol, and for which open source software already existed. Consequently we produced the final protocol mapping based on the Liberty Alliance recommendations, and published this to the user group. We then set about the implementation.

The methodology used for the implementation itself was based on prototyping for the user interface and the waterfall method for the core processing code. Several versions of the user interface were built and experimented with by members of the ISSRG until the final one was released to Glasgow for end user trials. The final task of the project (not completed at the time of writing) is to pilot the

implementation with a demonstration service built by NeSC, Glasgow. We will revise the interface if necessary as a result of the comments received, using funding from the EC TAS3 project.

Implementation

The project went broadly according to plan during the requirements gathering phase and conceptual modelling. The requirements were gathered by circulating the questionnaire to members of 12 international mailing lists (see Appendix 1). We received 26 replies and from these obtained 12 user requirements, some of which were mutually exclusive. We analysed the requirements against four possible attribute aggregation models, and found that none of them could fulfil all requirements, but identity linking was the best. It fulfilled 11 of the 12 requirements and required the lowest level of user interaction. The only requirement it could not fulfil was multi-hop proxying. Consequently we based our conceptual model on identity linking, and conceived of a new entity which we called the Linking Service. The conceptual model was produced and circulated to the user mailing lists for comments. All received comments were addressed and fed back into the conceptual model. This revised model was circulated again and all received comments addressed.

We presented our design and initial suggestions for protocol mappings at the Terena EMC2 meeting in Marseilles in Feb 2008 and at the Internet 2 Spring meeting in April 2008. There are always numerous ways that a message can be carried by underlying protocols and so in some ways determining "the best" protocol mapping was the most challenging part of the project. When we did finally get an agreement with our academic community users about the best protocol mapping to choose, this was frowned upon by the Liberty Alliance standard's group when we presented it to them in Stockholm in July 2008. They recommended an alternative mapping. So we then had to read a different set of LA specifications and work out the details of this new mapping. Once we had determined the new mapping and extensions that we needed (note that each of our potential mappings required some extensions to the standard protocol suite since attribute aggregation had never been a requirement of any of the standards groups' objectives), we finally set about the implementation in September 2008, only now we were five months late (month 19 instead of month 14). The lesson learnt is that when design is done by liaising with large groups of external people, then it takes a long period of time, usually longer than you had originally anticipated. Contingency planning is essential.

The implementation was challenging. The open source Liberty Alliance code had some bugs in it, and did not do everything we required of it. The original Discovery Service had several bugs particularly when processing schema and the resulting automatically generated code. This code relied upon an old version of Axis which produced unreachable code and never ending loops. These required minor changes to the original schema to ensure that runnable code was produced. There were also bugs in the database access code which meant that multiple instances of the same End Point Reference (EPR) object were generated and returned to the requesting service provider. The initial configuration of the service was also extremely complicated requiring the entire service to be configured at compilation time. Therefore changes were made to introduce configuration files that could be used to re-configure the service at runtime and reduce installation mistakes.

As the Shintau profile used a modified SOAP wsse Security header containing multiple security tokens rather than the traditional single token, support was added for the processing of this header. This ensured that the system could identify the user by a PID supplied by a linked entity and also verify that it trusted the provider of the initial act of authentication. The modified header structure also allowed for level of assurance support, based on the SAML 2.0 Level of Assurance draft [4], to be built into the system. This support took one of two forms depending on the installation type: If the Discovery Service was installed as part of a Linking Service (LS), then support was added to ensure that only accounts that were registered at a higher LoA than the session LoA could be returned to the requestor, preventing the user from creating links with low levels of assurance and using them at higher Session LoAs. If however the Discovery Service was installed as part of an IdP then support was added to ensure that only requests with a sufficiently high LoA could allow the user to access additional attributes.

In order to fully support the LS's security policies, link release policy (LRP) support needed to be added to the system. This was accomplished through the use of a new database table which stores

the LRP rules. A mechanism was implemented that checks these stored link release policy rules against the EPRs prior to the latter's release. This ensures that only EPRs that match the stored LRP rules are included in the response message.

The final change required to the discovery service's functionality was a mechanism that allowed the discovery service to map the identity information contained in the IDWSF Discovery Query request into an authentication session for the IdP. The IdP can then retrieve this information and use it to identify the user's linked account when the IdPs attribute authority endpoint is queried with an appropriate SAML 2.0 attribute query message.

The largest and most complicated change required to the LS was the addition of support for the signing and encryption of IDWSF Discovery Query messages and their associated responses. As the open source implementation utilised did not provide any support for these security features complete support had to be added using the Java Security API and the Bouncy Castle JCE. Related changes were also made to enable the service to use SAML 2.0 metadata at service initialisation for the creation of certificate trust stores and provider entries in the database.

Making the user interface user friendly was a particular challenge and very time consuming. The first prototype interface was far too technical. We simplified it several times during the course of the development in order to make it accessible to the average user. During this process, we encountered the following challenging problems:

- i) Displaying the Shibboleth name identifiers to users. These are completely user unfriendly, comprising of 128 bit numbers. In order to overcome this, we had to introduce a nickname feature. This allows the user to overwrite the 128 bit string with his own user friendly nickname.
- ii) The UK Federation meta data is deficient for our purposes since it does not contain a user friendly name for service providers. However, it does contain one for identity providers since this is needed for the Shibboleth Where Are You From Service. But we need an equivalent feature for Where Are You Going To column in the link release policy. Until this is added, the linking service will either have to be hand configured (tedious, error prone and time consuming) or it can be automatically configured from the meta data but it will then display user unfriendly names to the user. To date we have no satisfactory solution to this problem, since we require the UK Federation meta data schema to be updated.
- iii) A completely new vocabulary had to be chosen for the user interface, which did not mention any of the terms used in our conceptual and detailed design documents. Even mentioning "attribute aggregation" was judged to be too technical for the average user to comprehend and so had to be removed. Instead the interface talks about linking user accounts together. We hope our new jargon free vocabulary will be user friendly, but the proof of the pudding will be in the user trials to be carried out by Glasgow.

Once we had developed the software we demonstrated it live at the Internet 2 Spring Meeting in April 2009 and at the Terena Networking Conference in Malaga in June 2009. We also plan to demonstrate it at the EMC2 meeting in Rome on 22 October 2009. Finally we gave the software to NeSC Glasgow in June so that they could set up a pilot demonstration service and test it with end users.

Outputs and Results

This has been one of the most successful JISC projects that the ISSRG has undertaken. The outputs of this project have been innovative and of high quality, as well as being of practical benefit to users. Furthermore it has directly led to a significant amount of additional research funding being obtained.

The highlights of the project's outputs include:

- We are on the cover of May 2009's edition of IEEE Computer magazine in a special edition on Identity Management. The paper [1] describes the Shintau model and implementation.
- We have a detailed paper describing the Shintau conceptual model accepted for Next Generation Computer Systems journal [3]
- We presented a paper describing our user requirements survey at the E-Portfolio conference in Maastricht [2]

- We successfully demonstrated the software at the Internet2 Spring meeting and the Terena Networking Conference 2009 in Malaga.
- A live demonstration of the Shintau software is available for public use at <http://issrg-beta.cs.kent.ac.uk:8080/loademo.html>. This demonstration allows the user to create multiple accounts at up to 4 IdPs, then establish links between these accounts using the Shintau linking service, and finally to access various service providers that have different requirements for the number of attributes (from different issuing authorities) that must be presented in order to grant access. The user can therefore see for himself when he is granted access by aggregating his attribute accounts together and when he is denied access by not linking them together.
- We have secured further funding (€1M) to continue our research into attribute aggregation and federated identity management as part of the EC TAS3 Integrated Project which runs to end of Dec 2011. We plan to add support for attribute aggregation to Microsoft's CardSpace in the next version of our software.

Further outputs that are expected to arise directly from the Shintau project are:

- The RA on this project, George Inman, has been using the Shintau research for his PhD and he is now half way through, so we expect that Shintau will be primarily responsible for the award of a PhD in about two year's time.
- We will shortly release the Shintau software as open source code to the worldwide community, under the BSD license, after it has been validated by NeSC, Glasgow. This will allow sites with demanding security requirements to incorporate it into their applications
- We have prepared a full set of documentation for the Shintau software which is now being quality assured by NeSC. Once the documentation has been trailed by NeSC, its shortcomings highlighted and corrected by us, it will be released to the public along with the software

Outcomes and Impact

This project has been very successful from a security research perspective. It has been trailblazing in its results and outputs and in the design and implementation of a privacy preserving attribute aggregation service. It is believed to be the world's first public demonstration of a privacy preserving attribute aggregation service (i.e. the Shintau linking service) in which the user is totally in control of which attribute accounts he links together. But of course this means that the project's outputs are not yet being used by end users. Before this can occur, application integrators and developers will need to design and build attribute aggregation into their systems. It is functionality they are already demanding. We anticipate that it will be another 5 years or so before attribute aggregation is rolled out to significant numbers of end users (considering how long it has taken for Shibboleth to be rolled out into an operational service).

One significant short term impact is that the University of Kent has seen how Shintau's account linking service can be used in a proposed service that has been on its wish list for several years – that of providing students and alumni with a login account for life. A new proposal to the JISC AIM call proposes to integrate the Shintau results into a pilot Logins for Life service at the University of Kent, which will give students, alumni and staff a login account at Kent for their entire lifetime, thereby facilitate lifelong learning and stronger links with its alumni. Using account linking, users will be able to register their existing stranger accounts, such as with hotmail or google apps, with their Kent account, in order to access Kent's services.

Whilst the Shintau results are likely to have a large impact in the medium term, it is essential that the following issues are dealt with by the academic community in the short term:

1. The Shintau model of user interaction is similar in concept to the Shibboleth model in which users establish their own Attribute Release Policies (ARPs) prior to engaging in the use of federated services. To date, few if any UK Federation sites allow their users to establish ARPs. Consequently users are still not familiar with this concept and have no experience of using ARPs. We therefore do not know what the users' perceptions or level of user acceptance of ARPs will be.

2. With Shintau, the user has to establish his own Account Release Policy at the Linking Service before any of his linked accounts will be released to the federated services. Account Release Policies operates at a different level of granularity to ARPs. Until user's become familiar with ARPs they are quite likely to find the concept of Account Release Policies quite difficult to grasp. However we anticipate that once users do become familiar with ARPs, then understanding, setting and using Account Release Policies will be a relatively easy next step for them to take.
3. The Shintau software requires minor modifications to the use of existing Shibboleth protocols, therefore until these changes are integrated into the existing Shibboleth and SAMLv2 software implementations, neither IdPs or SPs will be able to use it
4. The Shibboleth and UK Federation MetaData is an impediment to the ease of use of the Shintau software. It needs to have a user friendly display name for SPs added to it, in the same way that it already contains a similar field for IdPs.

Conclusions & Recommendations

As a result of the Shintau project we make the following recommendations to JISC.

1. It is recommended that all UK Identity Providers start to engage with their end users, and empower them, by giving them more control over the release of their attributes to service providers. In other words, IdPs need to implement ARPs, and give user's full control over their contents. This will have a double benefit of
 - i) giving service providers the ability to apply finer grained access controls
 - ii) allowing IdPs to more easily conform to data protection legislation by allowing users to give their consent to the release of their attributes
2. JISC should pursue the inclusion of a user friendly display name for SPs to be included in its Federation Meta Data so that end users have an easily recognisable name to attach to the federation's SPs.
3. The Shintau conceptual model has several different modes of operation defined within it. During the life of the Shintau project we were only able to implement one of these modes, namely attribute aggregation by the SP from statically linked user accounts. Other designed but not implemented modes include:
 - attribute aggregation by the linking service. The advantage of this mode is that we believe that the SP does not require any special software to benefit from attribute aggregation, as all the hard work is done by the linking service. This option is well worth implementing and piloting as it may significantly reduce the overall cost of attribute aggregation,
 - dynamic linking of attributes. The advantage of this mode is that the user does not need to pre-establish his account links at the linking service prior to service invocation. The linking can be performed dynamically during service invocation. This option is well worth implementing and piloting since it may significantly increase user acceptance of attribute aggregation as it become more immediate i.e. when the user invokes a federated service he links his accounts at the same time.
4. The University of Brighton runs the Brighton and Sussex Medical School in conjunction with the University of Sussex. Students of the medical school are members of both institutions, and should be able to access resources with either set of credentials, but at present they can't do so because service providers don't let them. If Brighton and Sussex could automatically create links between their user accounts using the Shintau Account Linking Service, these students would never have to set up the links themselves and could then log in to the service provider sites from either of their accounts. As collaborative courses and other joint ventures become more common, it is likely that automated account linking will have a fairly wide range of applications. It would therefore be beneficial for JISC to provide a new appropriately secured web services management interface to the Account Linking Service which will let a management client, prompted as part of the student enrolment process, to create the account links automatically as soon as the student exists in multiple systems. This will save the student the effort of having to do this him/her self.

In conclusion, this project has been very successful from an e-infrastructure security research perspective. It has been truly trailblazing in its results and outputs. Attribute aggregation is currently a hot topic that is actively being discussed by the academic networking community, and we have been the first to produce a fully working demonstration and open source software. Of course, being at the leading edge means that the project's outputs are not yet being used by end users. However we anticipate that in another 5 years or so attribute aggregation will be widely rolled out to the end user community (considering how long it has taken for Shibboleth to be rolled out into operational services, this is not a long period of time). We believe that Shintau forms a solid rock on which to build future attribute aggregation efforts.

Implications for the future

We believe that Shintau forms a solid rock on which to build future attribute aggregation efforts. Since all the projects deliverables are in the public domain, and the software is open source under a BSD like license, then other researchers are free to take our software and adapt them to their own uses.

We currently do not know what the impact of Information Cards and Microsoft's CardSpace will be on the federated identity management arena. We are anticipating that it could be big, but one of its current limitations is its inability to aggregate multiple cards in a single transaction. For this reason our future research and development will be to look at how attribute aggregation can be built into the Information Card model and how the Shintau outputs can be utilised with CardSpace.

As stated above, the Shintau conceptual model has a number of different modes of operation (dynamic and static linking, SP and linking service aggregation) and we have only implemented one mode, namely static linking and SP aggregation. Other mode implementations are possible, and these will provide end users with different experiences and SPs with different amounts of implementation effort. These should be considered by JISC.

We currently do not know what the impact of IdPs offering different Levels of Assurance (LoA) will be. If IdPs start to implement LoAs, and SPs start to use them in their authorisation decision making, then this will have an impact on end users, who may wish to migrate to IdPs offering higher LoAs. The LoA offered by an IdP will impact on the user's experience of the Shintau linking service, since those IdPs offering low LoAs will inhibit the user from aggregating these attributes with those from IdPs offering higher LoAs. This may of course have a positive knock-on effect of encouraging IdPs to offer higher LoAs to their users. However, until LoAs are widely implemented, we really will not know what their impact will be.

Concerning the sustainability of the Shintau outputs, we are currently in a healthy position, having secured funding from the EC TAS3 project until December 2011, and we are continually bidding for more project funding. Therefore we will be able to continue to support and improve the Shintau software until at least this date. Furthermore, because the SP attribute aggregation functionality has been integrated with the PERMIS software suite (specifically with the Credential Validation Service) then as long as the PERMIS software is supported and maintained, the Shintau software will continue to be so.

References

1. David W Chadwick, George Inman. "Attribute Aggregation in Federated Identity Management". IEEE Computer, May 2009, pp 46-53
2. George Inman, David Chadwick, Nate Klingenstein. "Authorisation using Attributes from Multiple Authorities – A Study of Requirements". Presented at HCSIT Summit - ePortfolio International Conference, 16-19 October 2007, Maastricht, The Netherlands.
3. David W. Chadwick, George Inman, Nate Klingenstein. "A Conceptual Model for Attribute Aggregation" Accepted for Future Generation Computer Systems, Aug 2009
4. OASIS. "Level of Assurance Authentication Context Profiles for SAML 2.0" Working Draft 01. 01 July 2008

5. OASIS "eXtensible Access Control Markup Language (XACML) Version 2.0" OASIS Standard, 1 Feb 2005
6. David Chadwick, GansenZhao, Sassa Otenko, Romain Laborde, Linying Su and Tuan Anh Nguyen. "PERMIS: a modular authorization infrastructure". *Concurrency And Computation: Practice And Experience*. Volume 20, Issue 11, Pages 1341-1357, 10 August 2008.
7. Alfieri, R., Cecchini, R., Ciaschini, V., Dell'Agnello, L., Frohner, A., Lorentey, K., Spataro, F., "From gridmap-file to VOMS: managing authorization in a Grid environment". *Future Generation Computer Systems*. Vol. 21, no. 4, pp. 549-558. Apr. 2005
8. Tom Barton, Jim Basney, Tim Freeman, Tom Scavo, Frank Siebenlist, Von Welch, Rachana Ananthakrishnan, Bill Baker, Kate Keahey. "Identity Federation and Attribute-based Authorization through the Globus Toolkit, Shibboleth, GridShib, and MyProxy". Presented at NIST PKI Workshop, April 2006. Available from <http://middleware.internet2.edu/pki06/proceedings/welch-idfederation.pdf>

Appendix 1 Requirement Questionnaire Recipients

Members of the Jericho Forum (<http://www.opengroup.org/jericho/>)
OGF OGSA Working Group list
OGF OGSA Authz WG list
Liberty Alliance group working on attributes
Sun's Identity and Access Management group
The XACML TC
JISC-MIDDLEWARE-DEVELOPMENT list
IDENTITY-PROJECT-PUBLIC JISC mailing list
Terena EMC2 mailing list
Shibboleth Dev list
gsmv@webapp.lab.ac.uab.edu
the OSIS list (osis-dev@netmesh.org)