Common non run-time simulation services – lessons from UK MOD research

Keith Ford
Thales UK
United Kingdom
Keith.ford@uk.thalesgroup.com

Jon Lloyd
Dstl, UK MOD
United Kingdom
jplloyd1@dstl.gov.uk

Simon Skinner
XPI Simulation
United Kingdom
Simon.Skinner@xpisimulation.com

ABSTRACT

The United Kingdom Ministry of Defence (UK MOD) has a vision for utilisation of a range of common simulation services and components which aim to reduce Modelling and Simulation (M&S) acquisition and support costs and promote re-use across the UK defence enterprise. The Simulation Composition and Representation of Natural and Physical Environments (SCORE) is a joint industry and Defence Science and Technology Laboratory (Dstl) research project which is providing advice and recommendations to de-risk this vision. Under SCORE and the related Architecture Interoperability and Management of Simulation (AIMS) project, a range of studies into providing services and components as a service are being undertaken, with links into the NATO Modelling and Simulation Group (MSG)-136 Modelling and Simulation as a Service (MSaaS) activity.

This paper provides an overview of work undertaken within the first 2 years of the SCORE project associated with de-risking the delivery of non-runtime simulation services. This includes research to understand how data, behaviours, scenarios and other simulation components can be most effectively stored, codified and accessed. Work to date has drawn upon approaches undertaken in adjacent areas of defence to modelling and simulation, such as geospatial data, to evaluate candidate registry and repository solutions. The study has also included the development of data models and frameworks which can most effectively reflect data storage and retrieval requirement. The paper will discuss development and experimentation of these capabilities as part of efforts to inform a potential UK MOD approach to delivering Modelling and Simulation as a Service.

ABOUT THE AUTHORS

KEITH FORD has worked in the simulation industry for over 34 years and is currently the R&T manager at Thales Training & Simulation. During this time, he has worked on display systems, control loading, motion systems (for which he obtained his doctorate) and in the field of synthetic environments. He has been technical lead on many national and international collaborative research programmes and participated in the development of SISO's Distributed Simulation Engineering & Exploitation Process (DSEEP).

Jon Lloyd is a Principal Engineer working in Modelling and Simulation in the UK Ministry of Defence (MOD) Defence Science and Technology Laboratory (Dstl) for the past 16 years. During this time Jon has previously been a technical specialist in the development of CBRN Synthetic Environments and interoperability standards. Jon currently leads the Live, Virtual and Constructive Simulation research project in Dstl, overseeing research undertaken by Dstl, Industry and Academia on the development of simulation architectures, content and emerging technology exploitation related to the use of synthetic environments for Defence purposes. Jon has chaired and participated in various International Research Collaboration activities as part of his role.

SIMON SKINNER has worked in the simulation industry for 25 years, and is a founder and the Managing Director of XPI Simulation Ltd (XPI). XPI is the lead contractor on the SCORE project which collaborates with 15 other UK industry and academic organizations to deliver research to UK MOD. He is the chair of the SISO Common Image Generator Interface (CIGI) product support group and Vice Chair of the Reuse and Interoperation of Environmental Data and Processes (RIEDP) product development group. Simon was the winner of an award for the 'Best paper from around the world' at I/ITSEC 2012.

Common non run-time simulation services – lessons from UK MOD research

Jon Lloyd
Dstl, UK MOD
United Kingdom
jplloyd1@dstl.gov.uk

Keith Ford
Thales UK
United Kingdom
Keith.ford@uk.thalesgroup.com

Simon Skinner
XPI Simulation
United Kingdom
Simon.Skinner@xpisimulation.com

INTRODUCTION AND RESEARCH CONTEXT

The use of heterogeneous distributed simulation systems is becoming more common and accepted by the military. Applications include training at higher formation levels (constructive), as well as virtual systems for tactical training. Distributed simulations are also used for Concept Development, Experimentation, Mission Rehearsal and Test and Evaluation. SISO developed the Distributed Simulation Engineering & Execution Process (DSEEP) to provide a standardized and recommended process for developing distributed simulations. Even after many years' experience of developing distributed simulation systems, a lot of time, money and resources is still required during the non-runtime DSEEP steps i.e. steps 1-5, 7. Therefore there is still a need to investigate ways of making the non-runtime activities more efficient and cost effective. There are many reasons why it still takes a lot of time and effort to develop distributed simulations, which include the complexity of distributed simulations increasing to meet more sophisticated operational needs; integration of complex heterogeneous systems (both simulation and operational); more complex interoperability requirements as the fidelity of simulations increases; requirement to ensure a 'Fair-Fight' using multiple heterogeneous simulations; and wasted time, effort and money in regenerating data or simulations that already exists but is not known about. This paper focuses on addressing the last bullet point i.e. promoting the reuse of simulation services, simulation data and simulation assets. For the purpose of this paper Simulation Assets, Data and Services will collectively be referred to as Simulation Resources. Simulation Assets are physical systems such as Full Mission Simulators or operational equipment e.g. simulators or Battle Management Systems. Simulation Data includes any data required by a simulation application or the design of a distributed simulation e.g. visual/terrain databases, weather data. Simulation Services are executable code that provides some functionality run as web services, applications in containers or applications in virtual machines. e.g., Weapons Effect service.

UK SIMULATION INITIATIVES

The motivation for this research is that over the years the MOD has paid for similar simulation assets many times over i.e the terrain database for the same training area has been procured separately many times; the same platform models have been implemented many times in many different constructive simulation systems. Although at least some of these simulation assets satisfy specific requirements and necessarily suitable for reuse, e.g. low resolution for a fast jet simulator or high resolution for an urban training system, there is still a lot of duplication. In these times of austerity, the MOD needs to make best use of the limited budgets it has available. One of the best ways of doing this when procuring simulations is to reuse what already exists. The concept is not new, and a prototype repository for storing simulation assets was produced by the UK led Euclid Research and Technology Project (RTP) 11.13 project in 2003 [1]. The vision promoted by Euclid RTP 11.13 was that the MOD's from each of the 13 nations participating in the project (and potentially industries from those nations) would host their own repository node and that these nodes would be interconnected. This would enable searches to be made not only of a nation's simulation assets but also of simulation assets that other nations were happy to share with its coalition partners.

In the UK MOD there is a drive to ensure that simulation capabilities are supported by common services and infrastructure to drive overall through life cost effectiveness. A Training Transformation project was undertaken by the UK MOD from 2011 to 2013 which helped to identify a Service Oriented Approach (SOA) to the future delivery of simulation [2]. This identified a set of common components and services that would enable the UK MOD to

move towards the provision of Modelling and Simulation as a Service (MSaaS). The UK MOD is also researching setting-up a centralised Modelling & Simulation capability focussing on enabling a more effective, cheaper and environmentally sustainable capability for Defence. An interim 'Defence Simulation Centre' (iDSC), has been set-up to inform the design of a permanent DSC by demonstrating the benefits of ensuring technical coherence for simulation across Defence and with Allies; providing M&S advice and support across the Defence community; exploring the provision of common M&S services; and providing access to common verified and validated M&S resources promoting reuse and interoperability across Defence (see DTEC Catalogue below).

MSaaS is a common theme across international research activities and individual nations. The UK is participating in the North Atlantic Treaty Organisation (NATO) M&S Group 136: Modelling and Simulation as a Service, which has 17 nations contributing to understanding and de-risking the approach [3]. In 2013, the UK Defence Science and Technology Laboratory (DSTL), part of the UK MOD, initiated two significant 4 year simulation research projects to address these issues. Architectures, Infrastructure and Management of Simulations (AIMS) and Simulation Composition and Representation of Natural and Physical Environments (SCORE). In the context of this paper, AIMS is focussing on the reuse of simulation services, whilst SCOREs focus is on the reuse of simulation data and assets.

The AIMS Programme

The vision statement for the AIMS programme is, "To enable the delivery of a single environment, such that users can create capability from modelling and simulation components and services which are inherently interoperable". AIMS covers a number of activities but the one relevant to this paper is developing the concept of providing MSaaS. Figure 1 shows the Modelling & Simulation as a Service Ecosystem defined by AIMS. The stakeholders are Ecosystem Owner: responsible for maintaining the Ecosystem by providing the maintenance functions; Resource Provider: produces and maintains services and data that can be exploited by other users; Simulation Developer: uses the simulation services for developing simulation capability; Simulation User: Indirectly uses services by using simulation capability. The MSaaS Ecosystem comprises three main functions; Discovery service, Service Oriented Architecture (SOA) and Administration. It should be noted that the Discovery service in AIMS is wider than that required by SOAs as it also includes Discovery of simulation data and assets by human operators as well as machine-machine interactions. AIMS has generated the concept of a 'Simulation as a Service Delivery Architecture' [4]. This comprises the following services and tools: Discovery – of simulation assets, data and services; Orchestration – integration of services; Deployment – to local infrastructure or Community clouds; and SaaSDA Runtime tools – manages and monitors services. AIMS is investigating the use of Virtual machines, Docker Containers and web services for providing simulation services. It assumes that services can be running 24/7 or downloaded and run as required. Due to the complexity of the Orchestration activity this is the least mature of the concepts. AIMS is focussing on extending SISO's Federation Engineering Agreements Template (FEAT), a standardised format for recording federation agreements in XML to increase their usability and reuse, to capture the design of a distributed simulation. The output is known as a Simulation Composition, which can be stored for reuse. The Deployment activity generates a Simulation Deployment object, which can be used to automatically deploy a Simulation Composition to a specified infrastructure. AIMS has demonstrated the savings in time of being able to quickly deploy a previously designed distributed simulation so that it can be reused.

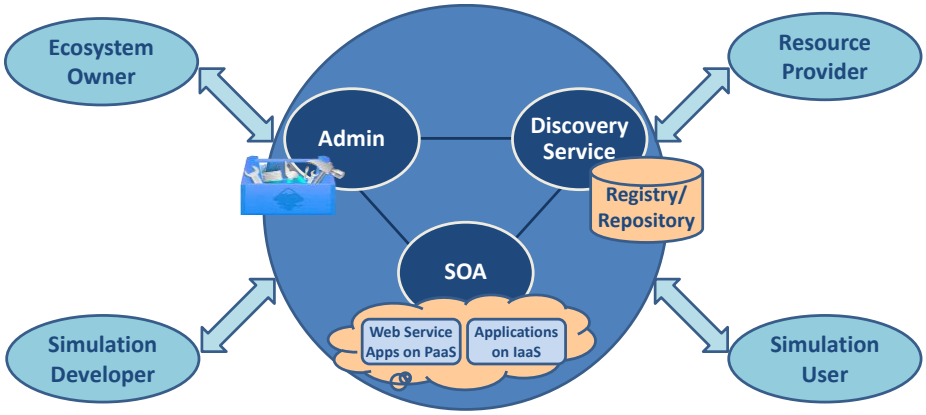

Figure 1: MSaaS Ecosystem

2.2. The SCORE Team

The vision statement for the SCORE team is, "To inform the development of coherent and consistent synthetic representations of the operating environment to provide more effective defence capabilities in support of training, concept development and experimentation, through-life acquisition and evaluation." Fundamental to the project is its mission to investigate and recommend methods to achieve a "Coherent and consistent representation of the operating environment for defence purposes". The SCORE project is researching next generation techniques to address both run-time and non-runtime simulation services. Run-time features include the components and services needed in a simulation, which might include environment and image generation services, along with terrain, weather and the action and reaction of computer generated forces. Non-runtime simulation services involve concepts such as the ability to store Data and Models (e.g. scenarios, orders of battle, terrains, platform models, 3D models, behaviours) in common repositories for re-use; reuse of Data and Models at run time within the common simulation components and services (to ensure that data and model knowledge of interchangeable components and services is retained for re-use); and methods of storing and retrieving expert knowledge and guidance – in the project this is known as the Guidance and Reference Information Package (GRIP). With respect to the reuse of simulation data, SCORE is focussing on the reuse of behaviours, scenarios, data (e.g., weather, databases) and components (figure 2). The driving force for this research is a desire to research the benefits of breaking down a CGF into functional components to enable a more flexible approach to providing the tactical environment in a distributed simulation.

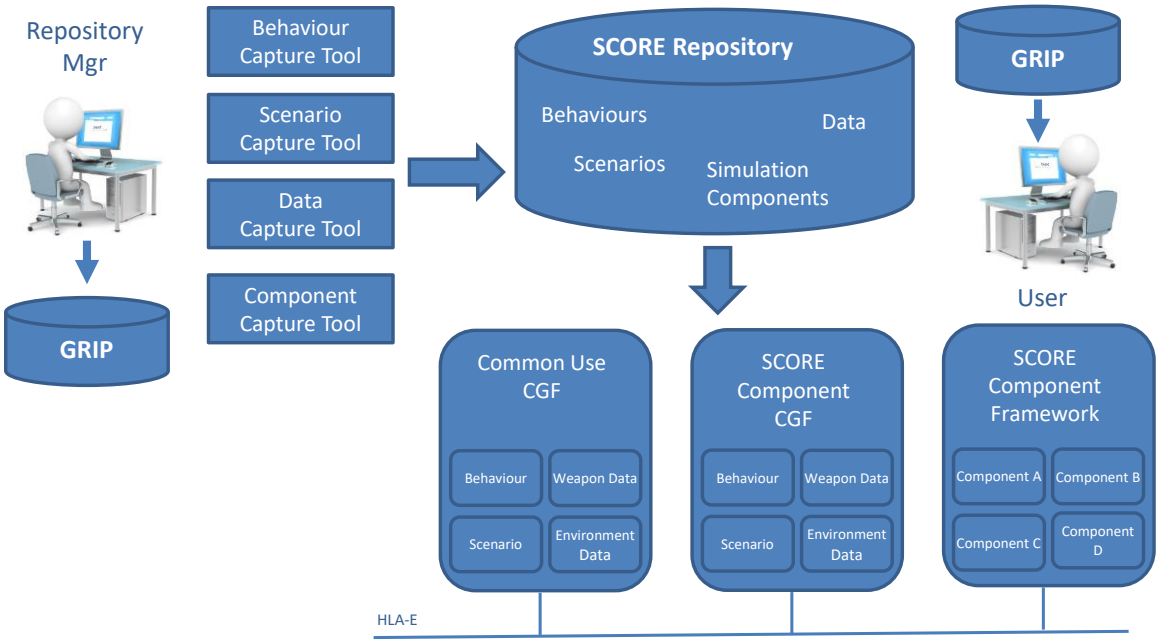

Figure 2: SCORE Overview

2.3. Other Simulation Reuse Initiatives

There are several international activities that overlap with the research being conducted by the UK. Probably the most mature is MSCO's Modelling & Simulation Catalogue and the Canadian Department of National defence (DND) Synthetic Environment Resource Repository (SERR) as they have fielded systems. The US also has the Enterprise Data Services (EDS) project (formerly known as Rapid Database Generation) [5], which is more aligned to the approach being adopted by AIMS/SCORE.

ISSUES WITH SIMULATION REUSE

Euclid RTP 11.13 produced a prototype repository, which was successfully demonstrated to key stakeholders from the 13 participating nations. Although everyone appreciated the benefits of promoting the reuse of simulation data and applications, when the project finished, there wasn't the will to exploit the concept. One of the reason why this may have been, is that it should be appreciated that reuse is not free. There is a cost to storing and maintaining

simulation assets over their lifespan. Also, making a simulation asset to support generic functionality when initially procured is much more expensive than if it was designed to meet a specific requirement. Unless the requirement for reuse is costed into the procurement, there is no incentive for a programme manager to take this hit. Unless future requirements are known in advance, it can be a costly and nugatory exercise to second guess subsequent functionality that may be required. There is a danger that the simulation asset is unnecessarily complex, which can introduce its own problems and costs. The most pragmatic approach for reuse is that simulation data and services are procured to meet their initial requirement and that new users will search for the simulation asset that most closely satisfies their requirements and make modifications to it. The modified simulation asset can then also be stored in the Repository for future use. Reducing the granularity of simulation data will help to promote reuse. As an example, SCORE is investigating storing vignettes in addition to complete scenarios and is developing a tool that concatenates vignettes specified in MSDL (Military Scenario Definition Language). As well as the cost of the modifications, there will also be the additional cost for revalidating the simulation asset to ensure that not only the new functionality satisfies the requirements but the modifications haven't altered the existing functionality in anyway. It should also be noted that even if a simulation is reused in an unmodified form, it may still need to be revalidated if used in a new situation.

DISCOVERY SERVICE

In keeping with the MSaaS concept, the ability for humans and machines to search for simulation resources will be provided as a Discovery service.

DSEEP Discovery Activities

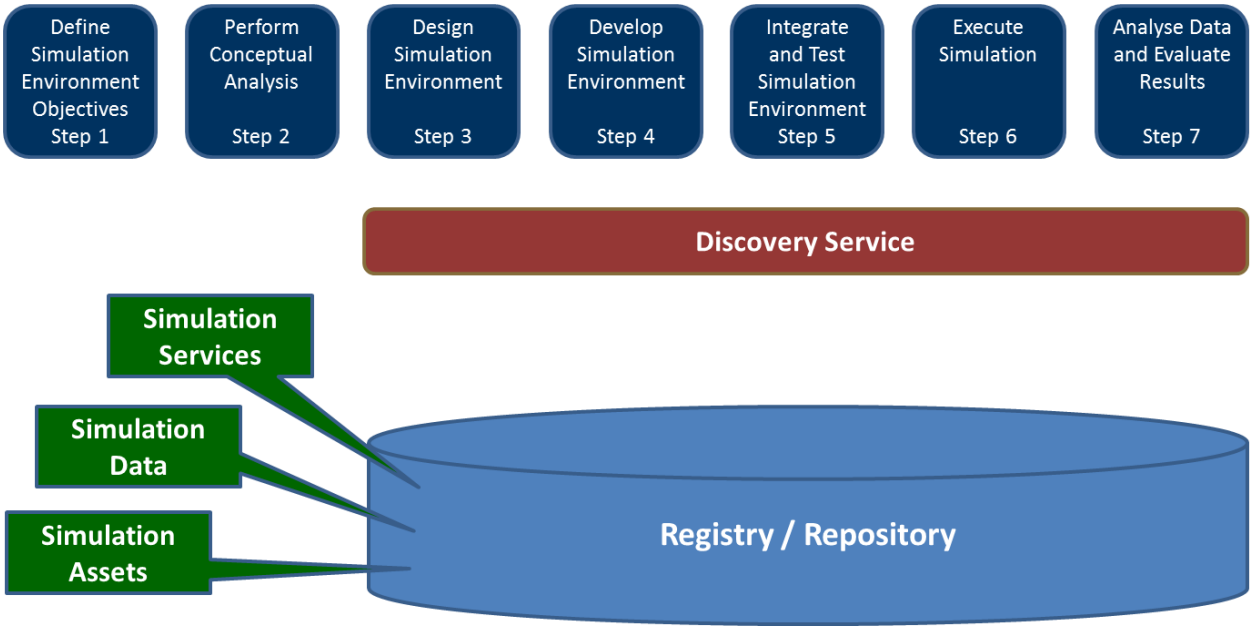

Figure 3: Discovery Service Support of DSEEP

The Discovery process underpins several DSEEP steps as shown in **Error! Reference source not found.**. In the Design Simulation Environment, Simulation Developers will search for existing simulation designs, services and simulation assets that satisfy the requirements of the conceptual model. In the Develop Simulation Environment, Simulation Developers will search for data e.g. visual database, required by a simulation or developer tools e.g. FEAT Editor to assist with developing the Simulation Environment. In the Integrate and Test Simulation Environment, Simulation Developer searches for tools to assist with integration and test activities e.g. HLA compliance tool. Machines search for and access/download services and data used by the simulation environment so that they can be deployed to the infrastructure. Execute Simulation is used if additional services have been included in the design of the simulation environment to provide resilience, these can be accessed/downloaded if the originals fail. For Analyse Data and Evaluate Results, searches can be made for analysis tools for processing the data logged at runtime. MSG-136 is developing an overlay to DSEEP for supporting MSaaS.

Discovery Concept

The concept of service or data Discovery is not new and is central for example to Service Oriented Architectures. SCORE is leveraging on the approach to Discovery used by the Geospatial community which is more mature than in the simulation domain. The Geospatial community has to manage petabytes of information stored in different locations, which is achieved using a combination of Registries and Repositories. The difference between a Registry and a Repository is that a Registry is like a catalogue that describes the data, whereas a Repository actually stores the data. The ability to reuse simulation assets requires them to be described using a well-defined and unambiguous data model. A core principle for Discovery services is to integrate diverse data centrally rather than mandate a single data model and template for all providers. Use of a Registry allows simulation assets to be described via a range of data and standards then mapped to the local template. To support MSaaS, the metadata needs to be readable by a human or a machine. Humans need to be able to read the information in the Registry in order to search for simulation assets, data or services and to read the metadata in order to decide what will satisfy their requirements. Machines will need to access the Discovery Service to access simulation services when a distributed simulation is being deployed to the infrastructure. The metadata will for example include information such as the capability of the simulation asset, data or service, the format in which data is stored, who is responsible for it and where it is located. For services, the Registry will contain a URL for accessing the service. It will also point to a repository entry where the executable code for the service is stored as a container or virtual machine. This enables Simulation Developers, rather than accessing the service directly at runtime, to download the service onto their own simulation networks so that they can be run locally. Physical simulation assets such as a flight simulator will not have a corresponding entry in the repository but the Registry will contain information for where it is located and contact details for using it.

The principle of the Discovery service implemented by AIMS/SCORE is shown in **Error! Reference source not found.** and is analogous to the concept of Service Oriented Architecture. The CSW-ebRIM Registry Service profile states that "A service-oriented architecture must support some fundamental interactions: publishing resource descriptions so that they are accessible to prospective users (publish); discovering resources of interest according to some set of search criteria (discover); and then interacting with the resource provider to access the desired resources (bind). Within such an architecture a registry service plays the essential role of matchmaker by providing publication and search functionality, thereby enabling a requester to dynamically discover and communicate with a suitable resource provider without requiring the requester to have advance knowledge about the provider" [6]. A Simulation Resource Provider stores the information in a Repository and its existence is published in the Registry. SCORE has extended the principle of the 'GetCapabilities' document used by GeoServer [7] by introducing a 'Simulation Artefact' document for describing simulation resources. A subset of the information in the Simulation Artefact document is used to generate the simulation metadata. Simulation Developers will evaluate the metadata returned by a search and will drill down into the Simulation Artefact if more information is required. If the information described in the Simulation Artefact satisfies their requirement, the user can access the resource/data using the URL provided or in the case of a simulation asset the contact details of the person responsible for it. As in principle it is possible to federate Registries, they can 'harvest' data from other Registries

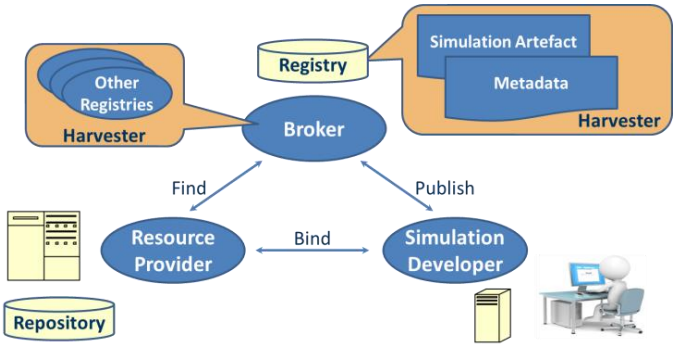

Figure 5: Simulation Resource Discovery Concept

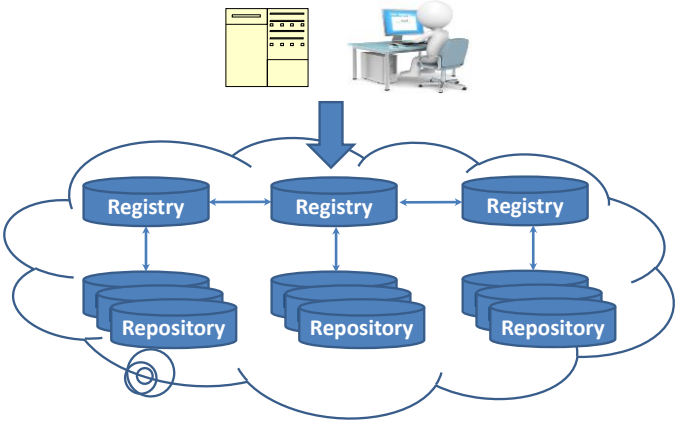

Figure 4: Multiple Registries/Repositories

provided by organisations that are willing to share their data. This enables a Simulation Developer/machine to Discover Simulation Resources described in multiple Registries from a single Portal. International standards exist that facilitates the harvesting of data from other compatible Registries. The Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) is a low-barrier mechanism for Registry interoperability. The SCORE concept enables a Registry to be integrated with many repositories. This enables Simulation Resource Providers to maintain their own Repositories, which provides them with complete control of their intellectual property. In some cases, the simulation resource may even be stored on a file system. From a Simulation Developer's perspective, they do not need to know or may not be aware of where the simulation resource is stored.

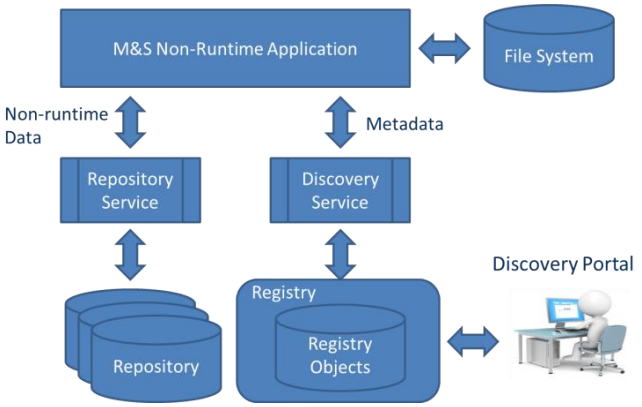

Figure 6: Integration of Registry/Repository with M&S Applications

Human access to the Registry is made through the Discovery Portal as shown in Figure 6. A Discovery service enables machine access to the Registry. As well as providing search capabilities the Discovery Service will also enable the creation and validation of metadata. Rather than M&S applications interfacing directly with the Repository, it is proposed that the application connects with a Repository service. This enables the application to only have to support one interface rather than having bespoke interfaces to different Repositories. Also, if the repository is moved, the URL only has to be changed in one place.

Discovery Service Vision

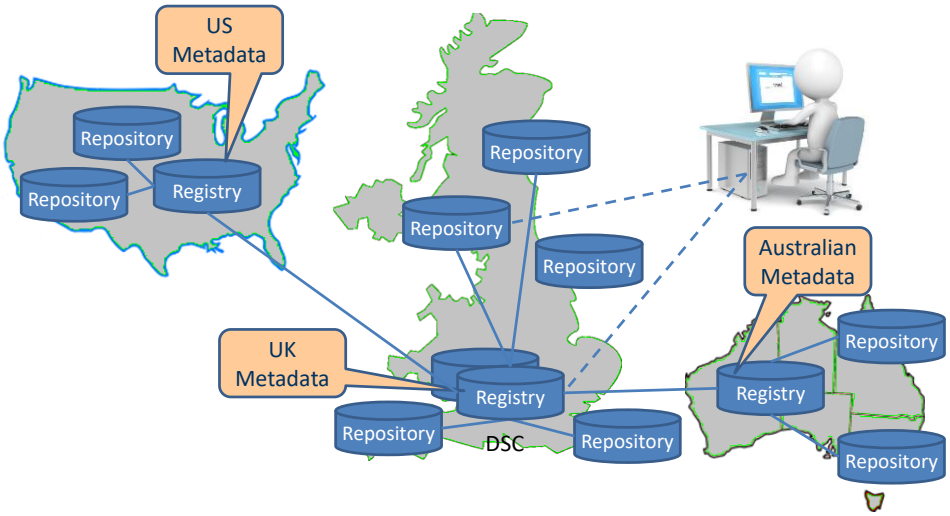

Figure 7: Discovery Service Vision

Figure 7 show a future vision for UK MOD providing a Discovery Service to facilitate the reuse of simulation assets. Although each organisation in MOD could manage its own Registry, this is considered an unnecessary overhead and it is recommended that a centralised capability is implemented, which can be accessed by MOD, coalition partners and potentially industry . It is envisaged that the Registry would be best managed by the proposed future Defence Simulation Centre (DSC). The metadata for describing the simulation assets in the UK should become part of the UK DTEC Modelling & Simulation Standards profile. The information held in the Registry would point to Repositories managed by different parts of MOD e.g. Army, Navy, Air force, Dstl. As well as MOD simulation assets being made available for reuse, other organisations such as industry and academia may also wish to promote their simulation assets as being available for reuse. These organisations may set-up their own Registries,

which may be federated with the DTESC Registry, or they may publicise their assets in the DTESC Registry. In this case, the Portal provided to a Simulation Developer will search not only for MOD simulation assets but also those from other organisations in the UK. Similarly, it is possible to federate Registries with the UK coalition partners e.g. NATO partners. The benefit of the approach proposed by SCORE is that it is not necessary to mandate a standard for the metadata. Different terms relating to the same parameter can be mapped enabling a Simulation Developer to perform searches across multiple Registries.

5. REUSE PROCESS

Figure 8 shows the activities that support the Discovery process.

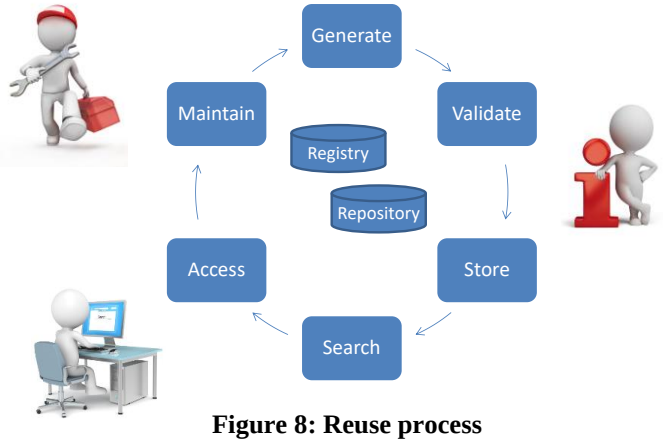

Figure 8: Reuse process

Generate. A simulation asset, data or service for reuse may be provided by a company that sells or licences a simulation resource (e.g., CGF, database, tool); simulation developer that has generated a simulation asset or data (e.g., simulation composition, database), software developer who has produced some executable code (e.g., component, simulation, tool); and Military Operator who provides military knowledge to be used by a simulation (e.g., ORBAT).

Validate. Before a Registry Administrator will accept a new Registry Entry, they will ensure that the metadata complies with the guidelines and that it accurately describes what is being offered. With respect to simulation data and services, additional testing maybe performed to confirm that the capability performs as described by the Provider. Tests may also be performed in a 'Sandbox' to verify that data or a service can be correctly accessed. Once the Simulation Asset, Data or Service has been approved, its existence can be published in the Registry.

Store. Repositories are only required for storing data or the executable code for running services. Being physical equipment, no information needs to be stored in the Repository for Simulation Assets most of the information required to access to them is contained in the metadata. Also, for services that are running 24/7, only a URL is required to point to them. As each Registry will have an internal Repository for storing metadata, this may also be used for storing small amounts of data that would otherwise be stored in an external repository.

Search. The simplest types of searches that can be performed are using key words. This is appropriate for some data but other types of data would benefit from more advanced techniques. As an example, searching for mapping data is best done by selecting a region on a map using a drawing tool, rather than providing the name of a country or region. There is a move in the IT industry to perform Semantic searches (as used by Google). This seeks to improve the accuracy of the search by understanding the searcher's intent and the contextual meaning of term. Most Registries/Repositories will have their own search engine and the type of searches that can be employed may influence which Registry/Repository will be used. The search capability will directly influence what metadata will be required to characterise the simulation assets. The ability to perform semantic searches will make accessing information about simulation assets much simpler for non-technical users. Searches may be performed for specific simulation assets that the Simulation Asset User knows exist e.g. OneSAF, or they may search for a particular capability e.g. Tornado GR4 simulations. The results from a search will generally return a number of options and the Simulation Developer will have to review the metadata for each one in order to determine which will best satisfy their requirement. A future vision is that the search could automatically be driven by the requirements for simulation platforms defined in a scenario. This will require translation of the MSDL describing the simulation platforms to the query language of the search engine being used.

Access. Simulation services running 24/7 will be accessed directly via a URL. Simulation Data stored in a repository (which in this context also includes executable code for running services) can be downloaded to the file system

using an appropriate repository tool. However, a better solution for reusing data is for the application that requires the simulation data to access from the repository via a data service.

Maintain. A Registry and Repository Administrator is required to maintain the information in the Registry/Repository. They will be responsible for the integrity of the information stored in the Registry/Repository by validating the metadata provided by Simulation Asset Suppliers. This to ensure that the metadata provided correctly describes the simulation resource to be stored in the Registry/Repository. The Registry and Repository Administrator(s) will also be responsible for retiring resources that are no longer appropriate and for managing access control. This is because some users may not be granted access to all the resources in the Registry/Repository.

6. COMPARISON OF REGISTRIES/REPOSITORIES

A number of Open Source and COTS Registries and Repositories have been reviewed by SCORE for implementing a prototype Discovery service. The criteria used for assessing the registries/repositories include 1) Availability/accessibility: Is the package a product or open source? If it is open source how easy is it to get access to? 2) Affordability: Is the product free with an open source licence or a commercial product? 3) Suitability: Which types of information can be stored? 4) Extensibility: Is it possible to extend the product or is it just a plug-in to other products? 5) Documentation: Is it easy to download the documentation? 6) Software requirements: Are there any specific requirements for operating the Registry/Repository? 7) Interfaces: What interfaces are provided with the tool to allow it to be integrated with other tools? 8) Standards: Which metadata standards does the tool conform to?

The MOD has recently let the Defence Geospatial Services (DGS) contract that has the objective to provide a consistent approach to MOD's management of geospatial data. As MOD has experience of using Envitia's GeoRegistry software on the DGS programme, it was appropriate to see if its capabilities could be extended to support the reuse of simulation resources. Before Envitia's GeoRegistry became an option, the SCORE team started evaluating the use of AtoM for providing search and storage functions. The investigation concluded that although AtoM's 'Discovery' capability wasn't anywhere near as advanced as Envitia's (as it is targeted at a different use), it could still be used as a repository for holding simulation data. It also has the ability to be federated with other AtoM repositories and uses the Open Archives Initiative (OAI) Protocol for Metadata Harvesting, which is also supported by Envitia's GeoRegistry. MSG-136 is using GitHub for storing Simulation Data and the AIMS/SCORE team intends to incorporate it in their architecture.

7. DATA MODELS

A Registry of simulation resources requires processing of unstructured, semi-structured and structured data and support queries that embody domain expertise. Data models that primarily use generalisation-specialisation (is a type of) or parent-child relationships (is a part of) are widely used, and is generally a default choice for organisation of data. This is logical, as classifications of this type impose a strong structure, enable use of a range of well-understood algorithms to navigate the data, and also mirror the way in which users of simulation systems tend to think about the problems that simulations solve and the information domains that they represent. The purpose of service metadata is to communicate to external entities the functions carried out by a service and the interfaces that it provides. A general purpose data model must enable services of all types to do this in a way that is meaningful to end users, whether these are human or machine, and which enables composition services to establish with complete certainty whether a service can interoperate with multiple other services. The approach taken in developing a service description template was to evaluate existing national and international standards and see how they could be extended to provide a broader capability. Applicable metadata standards that are relevant to describing simulation resources were identified and reviewed. These include Dublin Core Metadata Initiative (DCMI) – for networked resources [8]; ebRIM – e-business XML [9]; ISO 19115/19 – for Geographic information [10]; Web Ontology Language (OWL-S) [11]; and Resource Description Framework (RDF) [12]. ISO also has a standard that defines how a metadata registry should be structured. This is ISO/IEC 11179 [13] and defines the fundamental ideas of data elements, value domains, data element concepts, conceptual domains, and classification schemes essential to the building metadata standards. The main SOA standards all make use of architectural constructs and ontologies, which provides a framework in which the complexity of SOA-based systems can be described. This means that they all provide conceptual frameworks that enable flexible system composition and enable the description of multi-dimensional relationships between entities within systems and the concepts that they represent.

OWL-S is an ontology, within the OWL-based framework of the Semantic Web, for describing Semantic Web Services. It will enable users and software agents to automatically discover, invoke, compose, and monitor Web resources offering services, under specified constraints. The OWL-S ontology has three main parts: the service profile, the process model and the grounding. The service profile is used to describe what the service does. This information is primary meant for human reading, and includes the service name and description, limitations on applicability and quality of service, publisher and contact information. The process model describes how a client can interact with the service. This description includes the sets of inputs, outputs, pre-conditions and results of the service execution. The service grounding specifies the details that a client needs to interact with the service, as communication protocols, message formats, port numbers, etc. In addition, MSG-136 has evaluated the following metadata standards, which has also been taken account by SCORE [14]: NDMS (NATO Meta Discovery Specification) – supports the NATO Metadata Registry and Repository (NMRR) but is now obsolete; NCMS (NATO Core Metadata Specification) - replaces NDMS; UDDI (Universal Description Discovery and Integration); MSC-DMS (Modeling and Simulation Community of Interest Discover Metadata Specification); FEAT (Federation Engineering Agreements Template); and Swagger. The SCORE review of these standards concluded that ISO 19115/19 best met its needs. The work to date suggests that these standards can be used to discover simulation resources without adding any extensions and that only a controlled vocabulary is required.

8. IMPLEMENTATION

Envitia's GeoRegistry software has been deployed on a UKCloud (formerly Skyscape) Community cloud provided as part of the AIMS MSaaS infrastructure and is available to the AIMS and SCORE project teams for research purposes. GeoRegistry has been designed to process geospatial data. It therefore supports the ISO19119 standard for GeoServices and ISO19115 for GeoData and can process data from web services such as1) Web Features Service (WFS): provides an interface allowing requests for geographical features across the web using platform-independent calls; Web Mapping Service (WMS): protocol for serving (over the Internet) georeferenced map images, which a map server generates using data from a GIS database; Web Mapping Tiled Service (WMTS): protocol for serving pre-rendered georeferenced map tiles over the Internet; and Web Coverage Service (WCS): defines Web-based retrieval of coverages i.e. digital geospatial information representing space/time-varying phenomena. If the M&S data to be stored can be described and accessed using the above standards, then no changes are required to the GeoRegistry software. Although this may be appropriate for some data e.g. scenario data specified using MSDL, it will not be relevant to the vast majority of simulation data. In these cases, new metadata needs to be stored in the 'SimArtefact' document that describes what the resource is and how it is to be served. The areas requiring further development to process this additional data are highlighted in pink in Figure 9. The 'SimArtefact' documents are held in a local Repository to GeoRegistry. SCORE has also used the AtoM repository [15] for holding simulation data, which has been installed in the Skyscape cloud. As AtoM is open source, it provides a cheap and convenient way for organisations to store simulation data. Although the AtoM software doesn't natively support geospatial standards e.g. ISO 19115, it does support Dublin Core. This provides a basic capability for searching for simulation data. Future research could investigate how easy it is to extend the metadata used by AtoM and how to display other digital content not currently supported by AtoM.

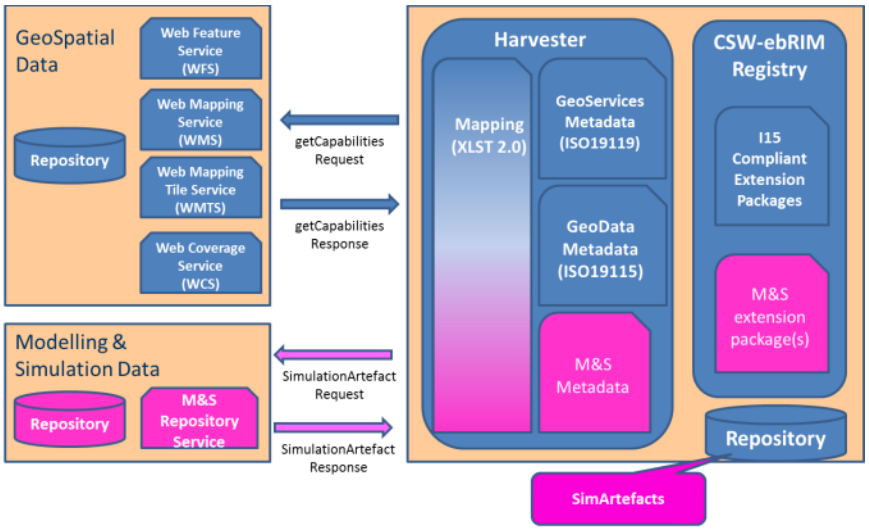

Figure 9: GeoRegistry Architecture

9. CONCLUSIONS

The AIMS Simulation as a Service Delivery Architecture was demonstrated in October 2015 and SCORE demonstrated its application to data reuse in March 2015. The Registry/Repository functionality was implemented using GeoRegistry and AtoM respectively. It should be noted that these applications were used as a convenient method to research the concept in the AIMS/SCORE projects and is not an endorsement for any future MOD capability. Leveraging of techniques from related disciplines e.g. geospatial data handling, is expediting the research and providing coherency. The use of a Registry capability, provides the ability to share information about simulation assets with Registries maintained by coalition partners. An advantage of this approach is that it is not necessary to mandate a metadata standard between nations. Each nation can define its own metadata, which can be mapped to comparable terms in other Registries. For example, 'color' or 'zip code' in a US Registry can be mapped to 'colour' and 'post code' respectively in a UK Registry. The SCORE team is keen to investigate harvesting information from other Registries as this would demonstrate the potential for interoperability with coalition partners.

One of the key drivers of this work is to understand how to enable machine-to-machine communication as a first step towards automating elements of the simulation management and composition processes. This is a difficult goal to achieve, and it must be considered from the start when developing data models. The advantages of breaking down capability into smaller parts i.e. to make a problem more tractable is true for all aspects of a simulation. Being able to describe and request very small and discrete elements of functionality should facilitate reuse. However, the ability to implement hierarchical services needs to be managed carefully to prevent unnecessary complexity and excessive latencies. Previous experience has shown that the problem with reusing simulation resources is not so much a technical problem but one of governance. It is one thing to set-up a Registry/Repository but another to ensure that its existence is publicised to all relevant stakeholders and that people use it. Discipline is required to ensure that the appropriate metadata is created when new content is added to the repository and that it is appropriately maintained. Ensuring that commercial constraints (e.g. IP ownership) are taken into account in developing technologies for a distributed repository and registry rather than purely a technical/engineering solution means that a usable system is being developed rather than a 'white elephant'.

It is felt that in these times of austerity the time is right to exploit this research as the reuse of simulation resources offers the possibility to make significant savings. However, for the concept to be successful, a champion in an influential role is required to agree to the initial investment for setting-up a reuse capability and to ensure that the appropriate governance is provided. In the UK, the implementation of the future Defence Simulation Centre (DSC) capability should provide the necessary oversight.

10. LESSONS LEARNT

We found that the approach used by the Geospatial community for Enterprise Content Management provides a useful starting point, but the simulation world is more complex. We haven't found a need to extend ISO

19115/19119 by adding new elements to handle simulation resources. Rather, we only have added a controlled vocabulary. Reuse is not free. There is a cost of creating and managing of metadata. We must design for reuse. It is important to quantify amount of simulation data/services to be managed to confirm whether approach is cost effective. The biggest risk for Enterprise Content Management failing is not having a high-level champion to promote the concept, counting for many previous failed initiatives. It is important to carefully consider the granularity of the data to be reused. Finally, implementation of approach requires a cross-domain security solution or air-gapped.

FUTURE WORK

The AIMS and SCORE projects are planning to further research the development of non-runtime services over the next 18 months of the contracts before making recommendations to the UK MOD. The research will continue to support NATO MSG-136 activities in understanding the operational concept and governance of such concepts as part of future MSaaS approaches, and will also participate in experimentation activities to de-risk the technical approach specifically providing the Registry/Repository capability as part of these efforts. Specific activities include investigating how to enable machine to machine communication to automate the discovery and integration of components into viable M&S applications; researching how to harvest information from other Registries through NATO MSG-136 to verify the mapping of different metadata; and linking the work carried out to date with run-time behaviours and automatic configuration of simulation services.

ACKNOWLEDGEMENTS

The AIMS and SCORE projects are funded by the UK MOD Chief Scientific Advisors research programme which is delivered by the Defence Science and Technology Laboratory (Dstl). The authors would like to acknowledge the contribution to this paper provided by the AIMS and SCORE teams.

REFERENCES

[1] E. Koster, M. Fabbri and R. Davies, “The Repository in Support of the Synthetic Environment Development & Exploitation Process,” ITEC 2003.

[2] S. Armstrong, “Training Transformation CCD Core Project Stages 1 & 2 Report,” NW/Pj/TTCCD/3981, April 2012.

[3] “NATO MSG-136 Terms of Reference,” 2014.

[4] K. Ford, “AIMS Simulation as a Service Delivery Architecture,” 2016.

[5] R. Scrudder, K. Gupton, B. Carlton, “Rapid Data Generation (RDG) – An Architecture Pattern for Simulation Data Sharing and Event Preparation,” SISO 14f-siw-054, 2014.

[6] “CSW-ebRIM Registry Service – Part 1: ebRIM profile of CSW (Catalogue Service for the Web), Open Geospatial Consortium,” 2009.

[7] “GeoServer,” [Online]. Available: http://geoserver.org/.

[8] “Dublin Core web site,” [Online]. Available: http://dublincore.org/.

[9] “ebXML,” [Online]. Available: http://www.ebxml.org/.

[10] “19115-1:2014 Geographic information -- Metadata -- Part 1: Fundamentals,” ISO.

[11] “OWL-S: Semantic Markup for Web Services,” W3C, [Online]. Available: https://www.w3.org/Submission/OWL-S/.

[12] “Resource Description Framework (RDF),” [Online]. Available: https://www.w3.org/RDF/.

[13] “ISO/IEC 11179, Information Technology -- Metadata registries,” [Online]. Available: http://metadata-standards.org/11179/ .

[14] K. Morse, M. Picollo, A. Bowers and M. Bertschik, “Developing Service Discovery Metadata to Support Modeling and Simulation as a Service,” 16f-siw-011, 2016.

[15] “Access To Memory,” [Online]. Available: https://www.accesstomemory.org/en/.