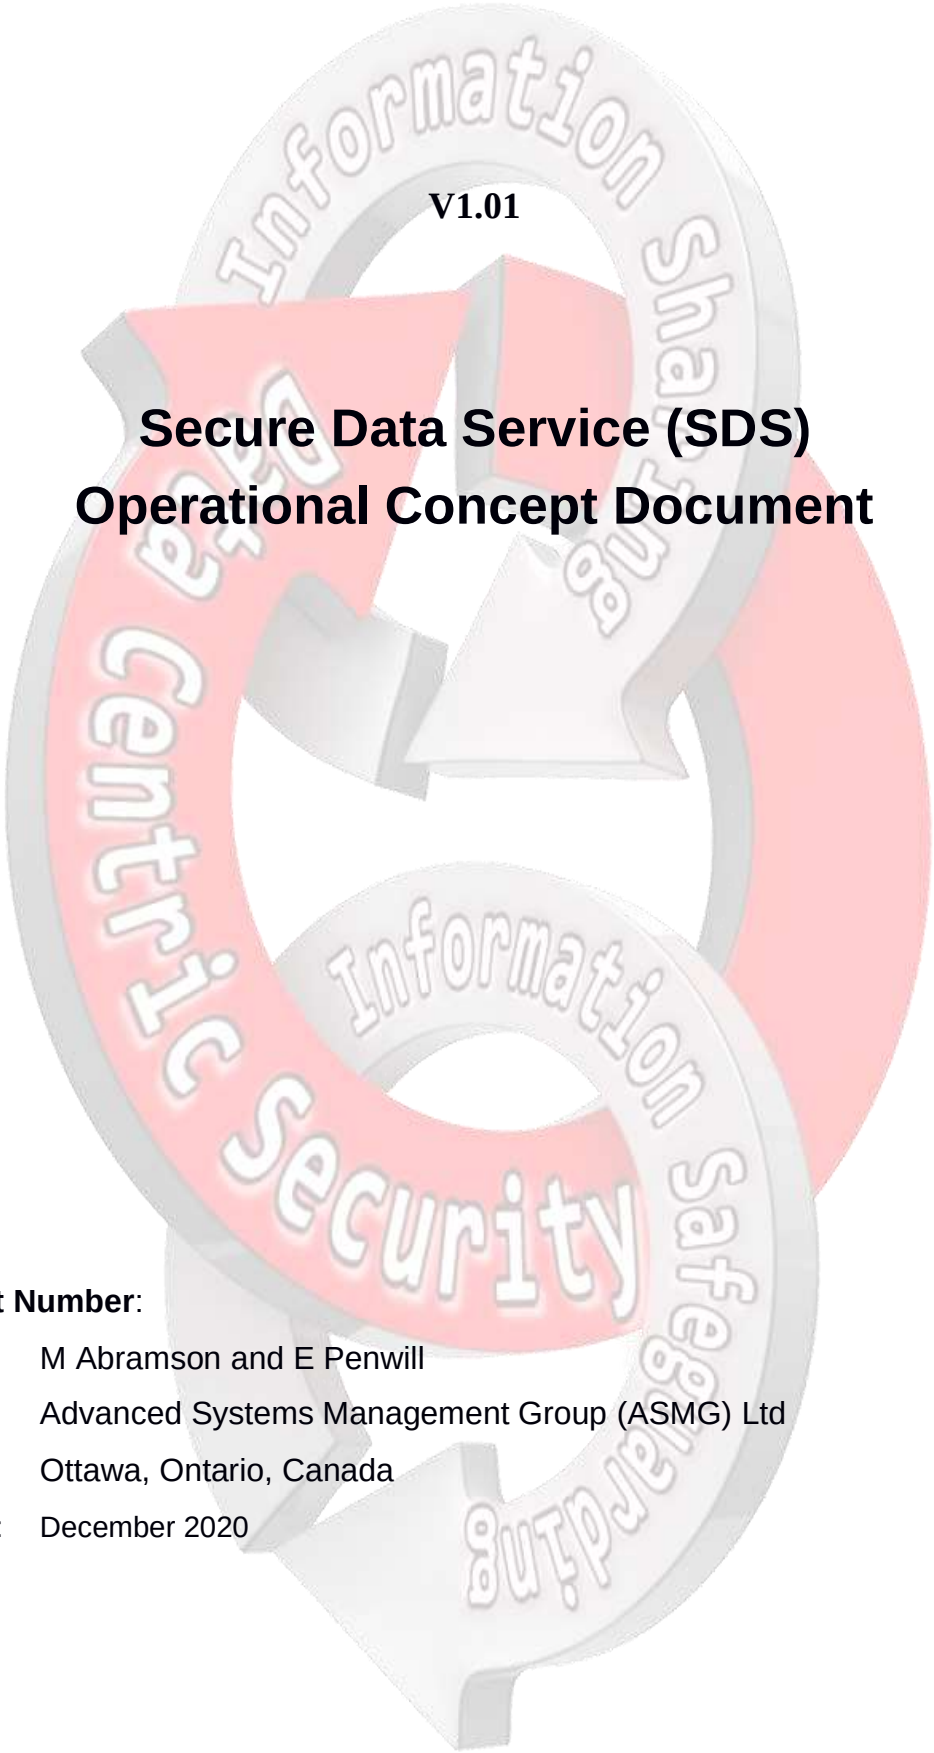

V1.01

Secure Data Service (SDS)
Operational Concept Document

Document Number:

Authors: M Abramson and E Penwill
Advanced Systems Management Group (ASMG) Ltd
Ottawa, Ontario, Canada

Issue Date: December 2020

Keywords

Interoperability, Standards, Information Sharing and Safeguarding, ISS, Data Centric Security, DCS, Secure Data Services, SDS, Information Exchange Framework, IEF™, IEF-RA™

Abstract

Operational Concept Document (OCD) for the Secure Data Service (SDS), describing the current interoperability situation, proposed solution, benefits, risks and required changes in terms of organisation, process, personnel and technology.

SUMMARY

This document outlines an operating concept for the implementation, operation and maintenance of an Information Sharing and Safeguarding (ISS) solution that can be securely deployed to multiple data and security domains and operating environments (e.g., on-Premises, Deployed Platform, Cloud and Hybrid). The solution delivers Data-as-a-Service (DaaS), including:

1. Data Centric Security;
2. Policy-Driven Data-Centric services for:
    a. Data Processing (parsing, transformation and marshalling) of received messages;
    b. Selective provision of data and information elements in accordance with the recipients' needs and authorisations, e.g.:
        i. Packaging (aggregation, transformation, labelling and redaction) of data for release;
        ii. Formatting and routing of information based on each recipient's information sharing agreement;
    c. Automated labelling of data and information elements, and messages; and
    d. Runtime administration of solution configurations and policies; and
3. Integration interfaces for:
    a. The users own security services;
    b. The users own cryptographic services;
    c. The users own system, information and security management services; and
    d. The users own trusted logging system.

The solution is defined and implemented to enable users to securely deploy sensitive information to mission environments or to exploit the cloud (e.g., IaaS, and PaaS).

This document is written from the perspective of military usage or deployment. There is nothing inherently military in the Secure Data Service (SDS) architecture, design or implementation that precludes its use in any public or private sector solution. The SDS represents an alternate configuration of services defined in an open international specification issued by the Object Management Group's (OMG) Information Exchange Framework Reference Architecture (IEF-RA; Reference H) and their alignment to conventional Cyber Security approaches and services.

- *This Page Intentionally left Blank* -

TABLE OF CONTENTS

Page

TABLE OF FIGURES

1. INTRODUCTION

1.1 AIM

With the evolving and growing desire of organizations to all-source exploit data to better inform decision-making processes, and create a decision advantage, there is a need for flexible, agile, adaptive, deployable and secure data containers that can be used to balance the need to provide and protect information. Containers that will provide each decision maker with the information they need and are authorized to receive. The containers should provide tools that enable users to:

1. Automate the enforcement of information security and protection policies;
2. Adapt policies during operation to provide the flexibility, agility and adaptability needed to achieve mission objectives;
3. Minimize threats to sensitive (i.e., private, confidential, legally-significant, and classified) data assets;
4. Share and safeguard multi-domain data and information in accordance with user defined policies (e.g., rules and constraints); and
5. Deploy as a stand-alone or embedded service for on-prem, cloud and hybrid environments.

Policy Driven: The adjudication and enforcement of rules and constraints derived from, and traceable to, user or community approved policy instruments (e.g., legislation, international agreements, regulations, directives, information sharing agreements, operating policy and operating procedures);

Data-Centric: The adjudication and enforcement of information sharing and guarding policies (rules and constraints) governing individual data and information elements; and

Information Sharing and Safeguarding (ISS): A set of capabilities that provide users with the ability to responsibly share information based on user needs, user authorizations and data sensitivity.

1.2 OVERVIEW

The Secure Data Service (SDS) outlined in this document describes the data container in terms of the All-Domain Command, Control, and Intelligence (AD-C2I) environment. The SDS defines an architecture pattern for a secure data (virtual) container. It describes an integration of the Information Exchange Framework (IEF[1]) and traditional security services to deliver Policy-driven Data Centric Security Information Sharing and Safeguarding within and across data domains. The SDS enables the secure deployment of data assets to any platform by wrapping data with services that:

1. Enforce user policies governing sharing and protection based on the data content; and
2. Utilize traditional security services (e.g., access controls, firewalls, secure operating systems and cryptography).

The SDS adds a layer of data protection to enhance the traditional application, platform, access control and network security services. An SDS does not replace traditional security and cyber services, it adds a layer of defence that integrates and builds on their features and functions.

[1] The Information Exchange Framework (IEF) and its Reference Architecture (IEF-RA) are trademarks of the Object Management Group (OMG) and described in the IEF Reference Architecture (https://www.omg.org/spec/IRF-RA), which is an open reference architecture for information Sharing and Safeguard (ISS) employing Data Centric Security (DCS) principles.

The SDS will significantly improve enterprise information sharing and safeguarding capabilities by enabling organizations:

1. To securely deploy data and information assets to multiple virtual platforms and environments;
2. To retain positive control over access to, and release of deployed asset to the individual data element level;
3. To leverage existing infrastructure investments; and
4. To exploit Data-as-a-Service.

The ability of organizations to maintain active control over access to, and release of their own data and information assets will increase trust of the data producers, owners and custodians, and foster their willingness to deploy and share information.

The SDS also provides organizations the ability to responsibly share information tailored to the recipients' specific needs and authorizations; enabling organizations:

1. To improve information quality; and
2. To maximize the availability of information to authorized recipients, while protecting sensitive data elements (i.e., private, confidential, legally-significant, and classified) against unauthorized access, release, use, expropriation, tampering or manipulation.

1.3 GOALS

The goals for this Operational Concept Document (OCD) include:

1) Describe the goals and objectives for the Secure Data Service;
2) Describe the SDS features and functions from an operational perspective;
3) Describe the SDS's impact on user and operator environments;
4) Describe how the solution can or should be used;
5) Facilitate understanding of the overall solution goals among users (including recipients of the products of the solution where applicable), buyers, implementers, architects, testers, and managers;
6) Form an overall basis for long-range operations planning and provide guidance for development of subsequent solution definition documents such as the solution specification and interface specification; and
7) Describe the user organization and mission from an integrated user/system point of view.

1.4 SCOPE

This OCD will be described in sufficient detail to allow stakeholders to understand the concepts and value provided by the SDS, without delving into technical details.

1.5 TARGET AUDIENCE

This OCD is provided to inform stakeholders (sponsors, architects, planners, developers, users, maintainers) about the scope, benefits and limitations of the proposed solution.

1.6 BACKGROUND DOCUMENTS

The following documents inform the definition and development of the Secure Data Service.

A: [C-M(2014)0016], Alliance C3 Strategy, dated 17 Mar 2014

B: [AC/322-D(2005)0053-REV2], NNEC Data Strategy, dated 14 Sep 2009

C: [AC/259-D(2013)0025-REV2-AS1, MULTIREF], Roadmap for Implementation of the Technological Aspects of the Connected Forces Initiative (CFI), dated 9 May 2014

D: [AC/322-D(2009)0046-REV1-FINAL], NATO Information Management Authority, NATO IM Strategic Plan, dated 18 Apr 2011

E: [MCM-0106-2014], NATO Federated Mission Networking Implementation Plan, Volume 1, Version 3.0, dated 14 Aug 2014

F: [TR/2020/SPW014853/xx], Standards Transformation Framework (STF) Operational Concept Document, dated xx August 2020

G: [TR/2020/SPW014853/xx], NATO Core Data Framework (NCDF) Operational Concept Document, dated xx August 2020

H: Information Exchange Framework Reference Architecture IEF-RA), October 2019, https://www.omg.org/spec/IEF-RA/

I: Information Exchange Packaging Policy Vocabulary (IEPPV), May 2015, https://www.omg.org/spec/IEPPV/

J: ENCLOSURE 1 TO IMSM-0149-2019 (INV), REVISION OF THE DATA CENTRIC SECURITY VISION AND STRATEGY PROPOSAL FOR THE ALLIANCE FEDERATION, INCLUDING THE NATO ENTERPRISE, 28 February 2019

K. C-M(2017)0062, NATO Enterprise Communications and Information Vision, dated 05 December 2017.

L. C-M(2015)0041-REV1, Alliance C3 Policy, dated 25 April 2016.

M. C-M(2007)0118, the NATO Information Management Policy, dated 11 December 2007.

N. C-M(2015)0003, NATO Federated Mission Networking Implementation Plan, dated 21 January 2015.

O. C-M(2008)0113(INV), The Primary Directive on Information Management, December 2008

P. C-M(2002)49, Security within the North Atlantic Treaty Organisation (NATO), 17 June 2002 including Corrigenda 1-12.

Q. C-M(2002)60, The Management of Non-Classified NATO Information, 23 July 2002

R. ADatP-4774 CONFIDENTIALITY METADATA LABEL SYNTAX Edition A Version 1 DECEDMBER 2017

S. ADatP-4778 METADATA BINDING MECHANISM Edition A Version 1 OCTOBER 2018

T. AC/322-D(2014)0010-FINAL, ANNEX 1, NATO CORE METADATA SPECIFICATION (NCMS), References: (a) AC/322-D(2014)0010-AS1, dated 14 January 2015 (b) AC/322-N(2015)0006-AS1, dated 3 February 2015

U. Allied Data Publication 34 (ADatP-34(K)) NATO Interoperability Standards and Profiles Volume 1 Introduction (Version 11) 3 Aug 2018

V. CF C4ISR Capability Development Strategy, July 2009

W. Defence CIO and CAF J6 Direction and Guidance, July 2020

X. Defence CIO and CAF J6 Direction and Guidance 2020, IM and IT, and CAF Joint C2IS Planning, Joint C2IS Interoperability, June 2020

2. CURRENT SITUATION

2.1 OVERVIEW

There is an existing and growing operational need to get "the right information, available to the right person at the right time in the right form to enable effective decision making" (Reference D). Stakeholders (e.g., decision makers) are seeking to employ all-domain (all-source) data and analytics to promote information sharing and decision advantage. However, there are counterbalancing requirements from data owners to protect sensitive information from unauthorized access, release, use, expropriation, tampering or manipulation.

The need to simultaneously address these counterbalancing requirements, led to definition, experimentation and testing of ISS and DCS capabilities that demonstrate the ability to:

1. Enhance the quality of the information available to decision makers;
2. Enable the secure deployment of information and data assets to on-prem, deployed, cloud, and hybrid environments;
3. Reduce the complexity and scale of IM/IT resources required for each mission or operational deployment;
4. Enable data producers, data owners, and data custodians to control access to, and/or release of, data and information assets;
5. Improve ISS and DCS monitoring and auditing;
6. Increase the users' ability to adapt ISS and DCS operations to changes in the mission and/or operational environments;
7. Minimize potential vectors of attack on data and information assets;
8. Provide Data-as-a-Service (DaaS);
9. Accelerate the data and information lifecycles; and
10. Deliver a Day-0 Capability.

More generally, provide the ability to responsibly share information, or "*the ability to maximize the availability of quality-information to decision-makers, while simultaneously protecting sensitive information from unauthorized access, release, use, expropriation, tampering or manipulation.*".

2.2 CAPABILITY SHORTFALL

At present, most IM/IT solutions are focussed on the sharing and exploitation aspects of the ISS equation. At increasing risk is privacy, confidential and security. What is needed is a service, or set of services, that enable users to achieve their own defined balance between:

1. Information sharing and safeguarding;
2. Data exploitation and protection;
3. On-prem or deployed (e.g., cloud, coalition networks, data lakes) environments; and

4. Allow users to implement solutions that enforce, and demonstrate conformance to, legislation, regulations, policy, Memorandum of Understanding (MoU), Information Sharing Agreements (ISA) and other ISS requirements and restrictions.

In other words, the ability to responsibly share information with users, allies and partners, or, maximise data availability to authorized users when it is needed, while simultaneously protecting sensitive (private, confidential, legally significant and classified) data from malicious or unintentional, and unauthorized access, release, use, expropriation, tampering or manipulation.

2.3 EVOLVING INTEROPERABILITY OBJECTIVES

The following sections briefly describe the current objectives for the items listed Section 2.1.

2.3.1 Increasing Information Quality

Decision makers at all levels, and across all sectors, are seeking to increase the availability and quality of the information available to the decision-making process. In a military context, this is a monumental task. As illustrated in Figure 1, military All-Domain Consultation, Command, Control, Communications and Intelligence (AD-C4I) involves multiple levels and layers of complexity in the collection, storage, processing (integration, fusion and curation), analysis, sharing and visualization of data and information elements. AD-C4I requires the integration of multiple data domains, employing a wide range of formats, vocabularies, syntaxes and semantics, the analysis of this data, and the production of quality decision ready information, in accordance with each users' needs and authorizations. The information must be bound with metadata to enable the information to be stored, discovered and safeguarded by software services that rely on metadata (labels) to execute their function. Historically, this labelling process was enabled through manual interventions, however, the variety, velocity and volume of modern information operations makes this manual intervention impractical. The information services need the ability to label data (e.g., STANAG 4774) and information artifacts based on user defined policies (rules and constraints).

Further fuelling this challenge is the scope of individual information requirements that often vary based on partner configuration, communications, mission threads, roles and responsibilities, and again by mission, operation, phase and command intent. In addition, individual decision makers may characterize information needs differently using qualities such as:

1. **Timeliness**: Received in time to render a decision and direct an action;
2. **Accuracy**: Free from error or defect; precise; exact;
3. **Relevancy**: Tailored to specific needs of the decision maker or decision;
4. **Completeness**: Provides all necessary and relevant data (where available) to facilitate a decision;
5. **Usability**: Presented in a common functional format, easily understood by the decision makers and their supporting applications;
6. **Actionability**: Capable of being acted on;
7. **Trustworthiness**: Accepted as authoritative by stakeholders, decision makers and users; and
8. **Integrity**: Protected from inadvertent or malicious release, modification, tampering or data loss.

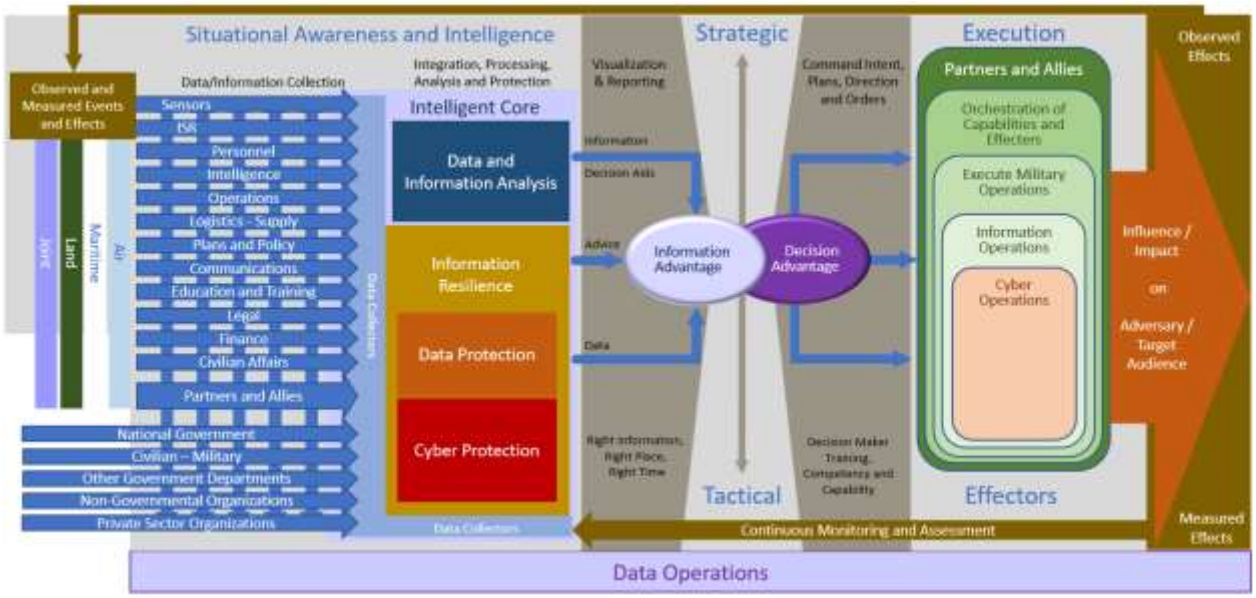

Figure 1: All Domain C2I (AD-C2I)

AD-C4I stakeholders are seeking solutions that provide methods and mechanisms that enable them to rapidly define, implement and certify ISS features and functions and deploy them to operations. These solutions must also provide them with the ability to adapt the operation of deployed solutions to the changing information needs of decision makers. Current ISS solutions are found to be rigid and brittle, therefore being unable to adapt to dynamic real-world situational and information needs.

Note: Many of the capabilities sought by the AD-C4I community have corresponding requirements in other domains including Finance, Healthcare, Personnel and more.

2.3.2 Deploying Capability

Stakeholders are seeking data storage and exchange capabilities that can be rapidly and securely configured and deployed to any platform, network and communication infrastructure, including:

1. On-Premises deployments such as headquarters using equipment such as:
    a. Servers,
    b. Desktops / laptops, and
    c. Mobile devices;
2. Coalition / partner[2] networks and infrastructure;
3. Operational deployed infrastructure and platforms;
4. Partner, Coalition or Public Cloud;
5. Shared Coalition and/or partner Data Lakes; and
6. Leverage the Public Cloud.

[2] Partner: National Government, Other Government Department (OGD), International Partner, Ally, Non-Government Organization (NGO), Private Sector Organization.

As the number of deployed features, functions, and information-assets grow, and the deployment-targets expand, so do the potential threats, risks and vulnerabilities. The non-traditional deployments also mean that the ability to define a boundary for the enterprise is rapidly vanishing, and along with it, the effectiveness of many of the traditional boundary security services and appliances. Stakeholders are seeking to place security safeguards around the data (i.e., put the boundary around the asset) or, to enable Data Centric Security.

2.3.3 Reducing Deployed Infrastructure

Stakeholders are seeking to reduce the number of distinct (virtually or physically separated enclaves) platforms, infrastructure and networks employed to separate and secure data at different levels of sensitivity or security. The reasons are varied but often include:

1. The costs and risks:
    a. To design, implement and deploy separate network infrastructure and applications;
    b. To manage and administer multiple networks, platforms and infrastructures;
    c. To integrate and maintain multiple networks, platforms and infrastructures;
2. The inability to effectively, efficiently and responsibly share information across domains and infrastructures;
3. The inability to maintain domain authorization levels as data and information elements are processed, aggregated, fused and analysed within the domain or network;
4. The cost of training for each variation of network, platform and infrastructure; and
5. The cost to update complex environments to incorporate operational and technological change.

2.3.4 Active Control over the Data and Information Environment

Stakeholders are seeking to increase their operational control over mechanisms governing access, use, storage and release of data and information assets. Many solutions rely heavily on as-a-service software, platform, and infrastructure vendors to secure and protect their data assets. Unfortunately, when reading the fine print of many as-a-Service agreements, it is clear that vendors place the responsibility to protect data and information asset on the user (/data-owner).

With many stakeholders also seeking to exploit the operational benefits (e.g., cost benefits, elasticity and flexibility) of public and hybrid clouds, they are relying on internal staff to understand and manage complex, often opaque externally developed and maintained as-a-service software, platforms, and infrastructures. As the initial hype fades into reality, stakeholders are realizing that the potential benefits, come with their own operational and security risks. Along with the migration to cloud services, stakeholders are seeking new and enhanced mechanisms to:

1. Rapidly define, develop/acquire, test, certify, and deploy new ISS capabilities;
2. Rapidly define, develop/acquire, test, certify, and deploy new services to secure and protect data and information assets;
3. Monitor and audit the operation of ISS solutions; and
4. Manage and administer ISS services in operations.

2.3.5 Monitoring and Auditing

Stakeholders are seeking to increase trust amongst data owners and custodians, and increase their willingness to share their data and information assets. To meet this objective, software, platforms and infrastructures must provide the ability to monitor and audit the collection, processing, analysis, sharing and visualization of data elements within and across systems, domains and organizations throughout the data life-cycle:

1. **Design Auditing**: The ability to review and analyse information, data, and solution architectures, designs and implementations to assure that user, and data-owner needs are addressed and that information policies (e.g., security) are effectively applied and enforced;
2. **Real-time Monitoring**: The ability to review and analyse run-time access, use, and exchange of data elements within and across mission components (e.g., systems, applications, middleware and networks). Solutions must also enable the monitoring of policy (e.g., Security) enforcement and any changes in policy environment or configuration of a service;
3. **Alerting**: The ability to identify issues (e.g., security or performance) occurring within the operating solutions and report then to administrators;
4. **Forensic Auditing**: The ability to analyse the architectures, designs and/or operational logs to verify that components are operating properly, and effectively enforcing information sharing and safeguarding (e.g., security) policies appropriately:
    a. **Design**: The ability to review and audit evolving solution architectures and design as a prelude to certification and providing feedback to the architecture and design processes;
    b. **Operations**: The ability to verify one's own components and configurations post operations;
    c. **Systems**: The ability to validate and verify that solutions and configurations are operating effectively and are candidates for certification for operations;
    d. **Policies**: The ability to validate and verify that ISA and ISS rules, constraints and configurations meet sharing, safeguarding, monitoring and auditing requirements, enabling their deployment; and
    e. **Partners**: The ability to validate and verify that partner-solutions meet the terms of MOUs and/or Information Sharing Agreements.

2.3.6 Minimize Vectors of Attack

Stakeholders are seeking to leverage the growing number of Information Management and Technology options in the on-Prem, hybrid and public cloud domains. Inhibiting the adoption of cloud and hybrid capabilities are concerns with:

- The virtualization of platforms and infrastructures blurring any delineation of boundary points in the environment;
- The skills gap in many organizations with respect to cloud development and operations;
- The delegation of security operations to data centre security service providers;
- The loss of monitoring and auditing capabilities in the cloud; and
- The growth of security threats in the environment.

Each of these concerns indicate the increasing number of possible threats or vectors of attack in this new environment, or simply old threats requiring new solutions (e.g., Zero Trust and Data Centric Security).

With boundary security being the most commonly understood form of cyber defence and the delineation of boundaries blurring in the cloud (virtual environments), a growing number of stakeholders are seeking to provide boundary security around the data and information assets, and carry their security attributes with them throughout their life cycle: e.g., Data Centric Security (DCS). The application of DCS will enable users to tailor defences to the specific sensitivity of the data being protected, and redact, where necessary, to assure responsible information sharing.

2.3.7 Flexibility, Agility and Adaptability

Stakeholders are seeking solutions that can be adapted to changes in real-world environments: including changes in user (e.g., decision maker) needs and authorizations. Traditional information sharing solutions are based on Application Program Interfaces (API) that codify the rules, constraints and configuration governing information sharing and safeguarding. These APIs are typically coded manually following traditional project practices, or more recently Agile/DevOps, practices. These practices severely limit the users' ability to securely adapt to the dynamics of real-world operations. Formally specifying the needed changes and recoding one or more APIs, then certifying it for operations may take weeks or months depending on the employed practices. The rigid and brittle nature of API maintenance practices are not conducive to ISS solutions for dynamic real-world operations: often forcing operators to find innovative ways to enable capabilities in the field, and possibly compromising information assurance.

Alternately, information exchanges are provided through back-office documents (reports, spreadsheets, and presentation material) that must be manually (human-in-the-loop) reviewed to provide operational value. This imposes production, approval and review delays that detrimentally affect operational effectiveness and efficiency.

Many stakeholders are seeking new or enhanced System-to-System (S2S) ISS and DCS solutions – that automate manual processes (e.g., labelling and data redaction) within the ISS stream. Services that provide flexible, agile, adaptive and secure ISS for:

1. Request/Response capabilities;
2. Publish and Subscribe capabilities; and
3. Event-driven global-update capabilities.

2.3.8 Data-as-a-Service

Enhanced data management is an imperative for decision makers at all levels (e.g., strategic, operational and tactical) and the amount of data in circulation and storage is increasing daily. Forward-thinking stakeholders recognize the value of data and seek to leverage it in decision-making processes, but most fail to use it to its full potential. Unfortunately, most data exists in silos bound to specific applications, systems and enclaves – thereby limiting access and its effectiveness. Keys to leveraging data to its maximum effect, include:

1. **Break Down Data Silos**: Enable users (e.g., individuals, services, applications and systems) to access real-time data streams from anywhere in the world;
2. **Use Data to Achieve Greater Agility**: Focus on the collection of data and the processing of that data into relevant data or information streams. Subscribers access the streams they need (and are

authorized to access, when they need them. Generate new streams as they are needed by analytics, services, applications, or systems to support mission threads or a specific mission thread;

3. **Manipulate Data More Easily**: Big data itself isn't useful, it is obtuse, disorganized, and has little use to most decision makers. The value of data comes from the trends and insights gained from closely analysing it and presenting it to the decision maker (e.g., data-driven hindsight, insight and foresight). DaaS aims to ease these constraints by offering catered data streams tailored to client needs;
4. **Single Source of Truth**: Agility is key to modern operations. Commanders (/decision-makers) must be able to quickly shift gears and refocus on new threats, risks, operational-parameters and objectives issues as they come to light. When you have more control over the data decision-makers use, it's easier to gain actionable insights from that data and leverage it appropriately. Breaking down solos reduces the likelihood that one decision-maker will overlook a data source that another organization or partner controls;
5. **Go Forward, Backed by Informed Decisions**: Understanding how important data is (and should be) to strategic, operational and tactical decision-making processes, DaaS offers a means to streamline access to data more effectively and with more acuity;
6. **Responsible Access and Sharing**: Maximizing the availability of data to authorized users, while simultaneously protecting sensitive data from unauthorized access, release, use, expropriation, tampering or manipulation is simplified when data is maintained by a single service – vs multiple stove-piped services, applications, systems and enclaves;
7. **Monitoring and Auditing**: Logging and auditing - who, what, where, when -- is simplified when data is maintained by a single service – vs multiple stove-piped services, applications, systems and enclaves; and
8. **Authentications and authorizations**: Are easier to manage when data is maintained by a single service – vs multiple stove-piped services, applications, systems and enclaves.

Stakeholders are seeking to deploy data-as-a-service so that:

1. Data is stored once and used by many users (e.g., services, applications and systems);
2. Data can be securely deployed to user specified operating environments: on-prem, cloud, or hybrid environment;
3. Data is secured and protected in accordance with their specified access, release, usage and sharing policies;
4. Data can be stored in its native form and transformed to coalition exchange semantics, application semantics or semantic reference models (SRM) as and when needed;
5. Data from multiple domains can be stored in a single store and released based on a users' needs and authorizations; and
6. Data provision can be rapidly adapted to address users' needs and authorizations.

2.3.9 Data Operations

Stakeholders are seeking Data Operations (DataOps) practices, procedures and tools to accelerate the data and information lifecycles and the delivery of information advantage to operations - from capturing relevant data, through delivering timely and accurate operationally critical information to decision makers. In an interoperability context, data services need to provide secure event-driven global-updates of data tailored to the needs of analytics, decision aids, and applications used by decision makers at all levels of the organization and across mission deployments. They must also provide decision makers with the ability to discover and access new data sources offered within the environment.

Administrators are seeking tools that enable monitoring and auditing operations and adapt data operations to changes in the operational environment, e.g.:

1. Threats;
2. Risks;
3. Objectives;
4. Plans and Orders;
5. Resources and Equipment; and
6. Organization (e.g., Partners, Roles and Responsibilities).

2.3.10 Deliver a Day-0 Capability.

Day-Zero (or Day-0) refers to the same day an incident, event (e.g., emergency, crisis or attack (e.g., man-driven or natural)) or vulnerability (e.g., cyber) is discovered. A Day-0 capability is the set of services and/or resources that can be employed to address or mitigate the incident, event or vulnerability on the day of discovery. In many instances, organizations take days, weeks or months to mount an effective response.

In many instances, the ability to capture, process and securely share incident specific information with colleagues and partners is critical to the planning and execution of an effective response. Stakeholders are seeking ISS solutions that provide at least a partial day-0 capability and the flexibility, agility and adaptability to rapidly advance and enhance that capability.

2.4 USERS OR INVOLVED PERSONNEL

Current solutions have developed, emerged, or evolved with varying degrees of independence, largely based on local needs and/or specific operational requirements. When designed and implemented they did not, or could not (e.g., out of scope requirements), take account of the broader enterprise requirements for operational integration and information interoperability. This has resulted in these solutions operating as stovepipes, providing only limited levels of information sharing and safeguarding capability. Information products (e.g., plans, reports, spreadsheets, and system products) are often used to mitigate shortfalls in ISS capability, where products are manually labelled by operators and shared using email and file share solutions.

These manual and often single domain (e.g., operational and security) solutions make it difficult (e.g., high risk and cost) to establish coalition networks with partners at differing levels of trust, including:

- Coalition partners;
- Other Government Departments (OGD);
- Non-Government Organizations (NGO);
- Private Volunteer Organizations; and
- Private Sector Organizations.

Highly skilled low availability operators are required to deploy and install current solutions and configure them for operation into single domain networks and infrastructures. Other highly skilled users are required to appreciate the sensitivity of each releasable product and label it so that traditional access control mechanisms (e.g., PEPs, guards and gateways) can control their release.

Many organizations are implementing systems that automate the production of information products. The sharing of these products is impeded by the manual practices associated with the reviewing and labelling of the products. As the volume of these system generated products increases, operators cannot be expected to efficiently review, evaluate and label each information product: products produced in real-time and at machine speeds.

2.5 SUPPORT CONCEPT

Traditional interoperability solutions are maintained by often under-funded operations and maintenance (O&M) activities. The design, development and deployment of broad-based reusable capabilities such as ISS or DCS often fall outside the scope of these O&M activities, and fail to materialize. These domain-specific solutions require high levels of SME involvement during development and maintenance processes to design, implement, test and certify mostly "point-to-point" data exchange solutions. The result is, APIs that are often rigid and brittle, poorly documented, error prone, not reusable, and difficult to maintain – resulting in the need for increasing numbers of resources and increases in risk and cost. Because of the low-level (technical) nature of API development, governance and oversight is often omitted, further limiting broader levels of ISS to evolve.

Constant changes in the information need (content and format) by enterprise and mission partners results in significant numbers of code changes to the interface software and other aspects of the solution. These changes also require recertification of each change by over-stressed teams prior to operations. The risks and costs associated with the implementation, certification and deployment of code changes stress strained O&M budgets and resources.

Operational requirements to significantly increase solution capability often requires the creation of new projects or project (/contract) amendments, with a defined (limited) scope to secure funding and navigate the procurement processes. The projects, working with a fixed set of requirements have again inherent inability to take account of the broader requirements for operational integration and information interoperability. This results in ongoing and difficult coordination, often with associated backwards compatibility issues and reduction in mission capability.

Scaled Agile: A set of organization and workflow patterns intended to guide enterprises in scaling lean and agile practices to plan, prioritize and manage capability development. Scaled Agile enables an enterprise to expand Agile development practices beyond the application development process.

Agile Development: Practices approach discovering requirements and developing solutions through the collaborative effort of self-organizing and cross-functional teams and their customer(s)/end user(s). It advocates adaptive planning, evolutionary development, early delivery, and continual improvement, and it encourages flexible responses to change

DevOps: Practices that combines software development (*Dev*) and IT operations (*Ops*). It aims to shorten the system development life cycle and provide continuous delivery with high software quality.

DevSecOps: Integration of security evaluation and testing at every phase of the software lifecycle, from initial design through integration, testing, deployment, delivery and maintenance.

In many environments, deployed solutions are extended and/or reconfigured by skilled operators to address immediate mission needs, bypassing development practices and procedures. Though immediately useful, these solutions are rarely transferable to a permanent capability as institutional memory is rarely maintained after mission completion.

Many organizations are seeking to employ Scaled-Agile, Agile Development and DevSecOps to address these limitations in their capability delivery capacity.

3. SDS BACKGROUND, OBJECTIVES, AND SCOPE

The Secure Data Service (SDS) addresses many of the capability shortfalls identified in Section 2.2. It delivers a configurable set of services, based on the OMG IEF-RA (Reference H), that enforce and demonstrate conformance to legislated, regulatory, policy-based ISS requirements obligations and restrictions. The SDS delivers policy-driven Data-Centric Information Sharing and Safeguarding within and across data domains.

The SDS ingests a user defined policy environment (Grey elements in Figure 2) comprising:

1. Semantic Policies;
2. Information Exchange Specifications;
3. Business Rules and Decision Logic;
4. Transformation Library;
5. Message and Metadata Schemas;
6. Message Parser Library & Mapping Files; and
7. Message Publisher Library.

The SDS services adjudicate and enforce these policy elements to ensure that the SDS only provisions authorized content to each request for information based on the requestors needs and authorizations.

Policy Driven: The adjudication and enforcement of rules and constraints derived from and traceable to user or community approved policy instruments (e.g., legislation, international agreements, regulation, directives, information sharing agreements, operating policy, and operating procedures);

Data-Centric: The adjudication and enforcement of information sharing and guarding policies (rules and constraints) governing individual data and information elements; and

Information Sharing and Safeguarding (ISS): A set of capabilities that provide users with the ability to responsibly share information based on user needs, user authorizations and data sensitivity.

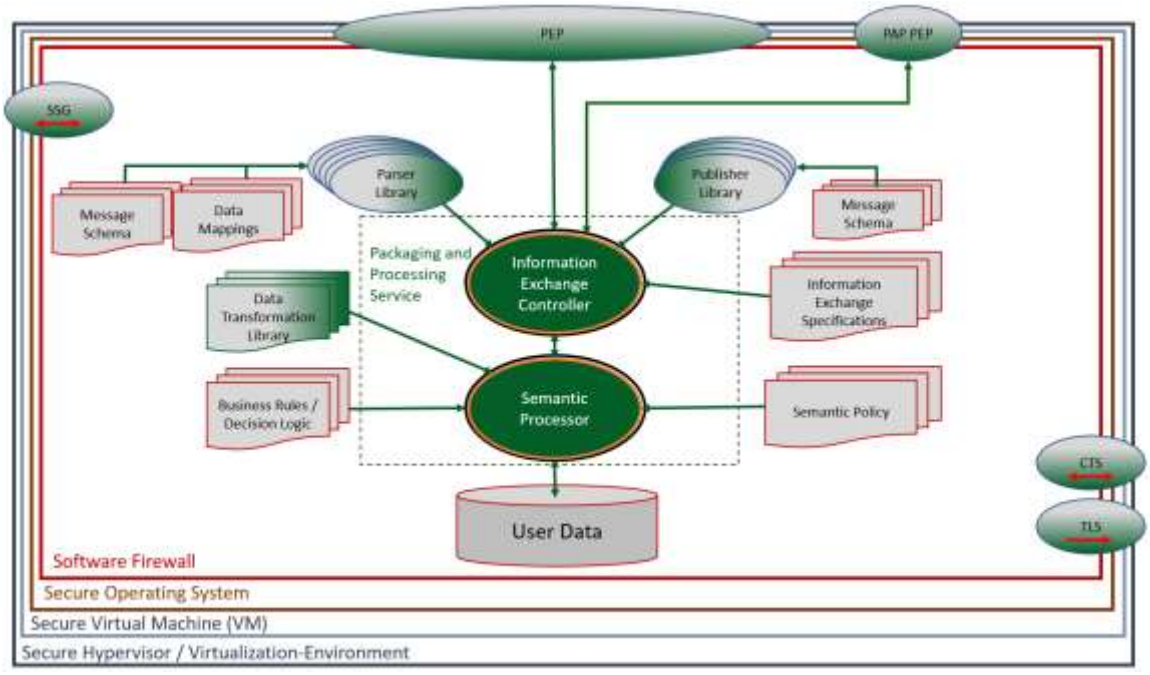

Figure 2: Secure Data Service

SDS services (Figure 2) include:

1. The IEF Services:
    a. Packaging and Processing Service (PPS): The PPS enforces semantic policies defining:
        i. The processing (parsing, transformation and marshalling) of received message content to the users specified data store(s);
        ii. The packaging (aggregation, transformation, labelling and redacting) of data and metadata elements to meet each recipients' needs and authorizations;
    b. Policy Enforcement Point (PEP): The PEP is the integration point between the users' infrastructure and the SDS services and enables the user or integrator to integrate access controls to the receipt and release of messages. The user or integrator would provide the Policy Decision Point and integration logic;
    c. Security Services Gateway (SSG): The SSG is a specialized PEP that provides peer-to-peer integration with the users cyber/security infrastructure (e.g., Identity, Credential and Access Management and Key Management);
    d. Trusted Logging Services (TLS): The TLS is a specialized PEP that provides peer-to-peer integration with the user specified logging system (e.g., Secure Database, or Distributed Object (e.g., Block Chain)); and
    e. Cryptographic Transformation Service (CTS): The CTS is a specialized PEP that provides peer-to-peer integration with the user specified cryptographic services.
2. Traditional security services provided by:
    a. Secure Operating System: The SOS (e.g., CentOS, and Open BSD) with its policies clamped to enable only the services needed for the OS and the SDS services; and
    b. Firewall: A traditional software firewall that clamps all the ports and protocols not required by the SDS and the Users' specified infrastructure (e.g., Middleware, and Security services).

Not illustrate in Figure 2 are the Policy Decision Point (PDP) and Access Control Policies typically associated with a Policy Enforcement Point. For the purposes of this operating concept, these elements are considered part of the Users Access Management Environment. The SDS PEPs provide the integration point through which the SDS is aligned and integrated with the users' environment.

3.1 TARGET CAPABILITY

Increasingly, decision-makers are seeking to exploit all sources of data to better inform the decision-making process. For this to occur data must be captured, stored, processed, fused, analysed, packaged and shared between multiple environments, organizations, systems, applications and individuals based on their need and authorizations. At all times, users, specifically data producers, owners and custodians want to assure that sensitive data is protected against unauthorized access, release, use, expropriation, tampering or manipulation. As illustrated in Figure 3, this exploitation requires the abilities to gather, store, process, classify, control, analyse and visualize data in ways that maximize the utility and effectiveness of that data to the decision maker. It is the ability to aggregate, integrate, fuse, analyse and visualize data from multiple sources in real-time that will ultimately improve the effectiveness, timeliness, and quality of today's operational decision-making processes and deliver real information and decision advantage. However, this must all be executed in accordance with the rules and constraints contained in data security policies and directives.

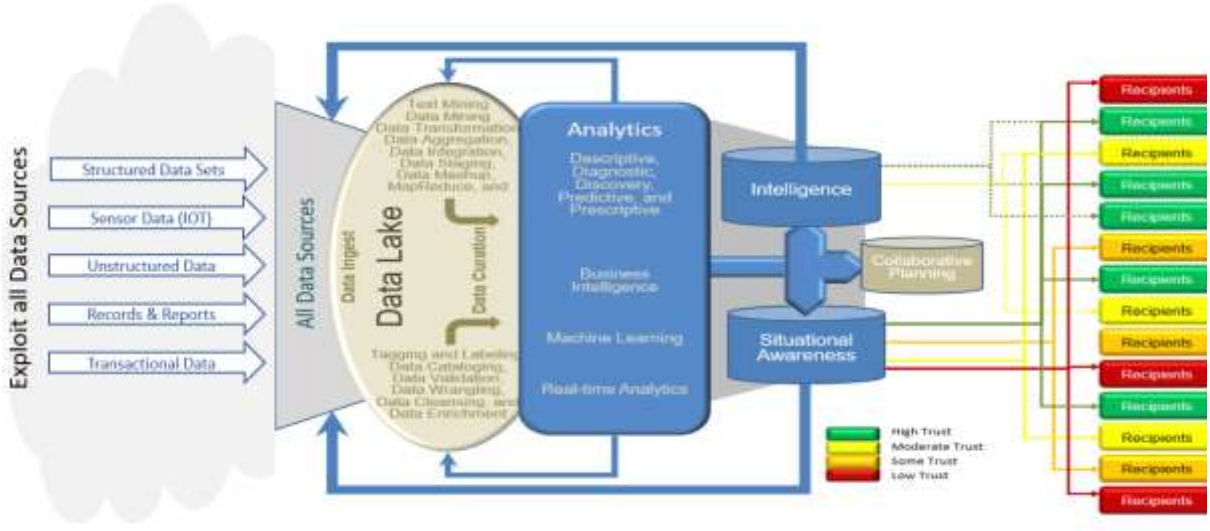

Figure 3: Data Exploitation

The SDS (Figure 2) provides a common data service that provides users with the ability to collect data (and/or metadata) from heterogeneous sources in either its native form, or normalized form (e.g., to a common Semantic Reference Model (SRM), Ontology or Schema). The SDS adjudicates and enforces user defined ISS policies that can be tailored for each mission environment, including:

- Participants needs and authorizations / Information Sharing Agreements (ISA);
- Desired Outcomes / Command Intent;
- Operational phase;
- Data sources;
- Exchange Infrastructure;
- Storage technology and Semantics/Schema;
- Labelling standards (e.g., STANAG 4774); and
- Binding standards (e.g., STANAG 4778).

ISS policies govern and orchestrate the processing and packaging of available data. Separating the policies (rules and constraints) and the software service that adjudicates and enforces those policies provides individual users with the ability:

1. To develop and deploy ISS policies tailored to mission and coalition configurations;
2. To ensure policies address both information needs, outlined in Information Sharing Agreements (ISA) and National Information Security policies and directives;
3. Automate the labelling of data based on the content of a specific information exchange, enabling other security services (e.g., PEPs, Guards and Gateways) to operate on the data;
4. To adapt policies to address changing mission requirements; and

5. Audit ISS against a documented baseline.

3.2 BACKGROUND INFORMATION

3.2.1 Information Exchange Framework

The SDS represents a new configuration to the Information Exchange Framework (IEF™) (Reference H) defined components. The IEF provides a reference architecture (Reference H) for DCS capabilities for:

A. Email;
B. Chat / Text Messaging;
C. File Sharing; and
D. System-to-System (S2S) Messaging.

The SDS represents a specific design pattern for the IEF focussed on S2S data and information exchange. One that will enable users to securely deploy system data to partner or coalition networks or the cloud. It combines the features of the IEF components with the security features provided by virtual machines, secure operating systems and firewalls. It builds on the IEF to define design patterns for a secure Data-as-a-Service (sDaaS) that users can deploy to multiple environments of interest to stakeholders: e.g.: on-Prem, deployed (e.g., deployed HQ, coalition networks, vehicles and mobile devices), cloud and/or hybrid environments.

3.2.2 Information Exchange Packaging Policy Vocabulary

The SDS services (i.e., Semantic Processor and Information Exchange Controller) employ the Information Exchange Packaging and Processing Vocabulary (IEPPV™: Reference I) for the design and implementation of ISS policies that can be translated to a runtime set of rules and constraints, ingested by the SDS, and used to adjudicate and enforce:

1. Information sharing agreements; and
2. Data security rules and constraints.

The IEPPV provides a light weigh ontology (e.g., vocabulary) and UML™ profile for defining ISS policies for processing, packaging and exchanging data elements.

3.2.3 DCS Vision and Strategy – In a NATO Context

The DCS Vision and Strategy (Reference J) supports a transitional process and emphasises an evolutionary path, through a set of defined DCS maturity levels, which consistently enhances the Alliance Federation's ability to achieve and maintain Information Advantage.

Within a DCS enabled environment, metadata (e.g., STANAG 4774) is used to describe and categorize data including the security classification of the data, identification of data ownership (and custodianship), retention and disposition, reusability and comprehensibility, and discoverability (Reference O). Together with Security Policy (Reference P) and the Management of Non-Classified NATO Information (Reference Q), it determines required security mechanisms for access control (including read, write and deletion), transmission, mediation between security domains within the NATO Enterprise, release beyond NATO CIS (NATO nations and non-

NATO entities) and usage of the information (including when shared). The SDS extends DCS by providing users with the ability to define policies that aggregate, transform and bind metadata to the content being packaged for release. This metadata is then bound to the messages for use by other DCS components (e.g., PEPs, Guards and Gateways).

DCS does not replace the existing security measures for confidentiality, integrity, nonrepudiation, authenticity and availability protection of the data and information environment. These existing security measures provide security of information at multiple layers, including boundary, network, endpoint and application [Reference K].

DCS has three main tenets where data is the key focus: Control, Protect and Share.

3.2.3.1 Control

Determine, based on originator-defined rules (/policy), where the data can reside and how it can transit across the Alliance Federation; provide a means for defining originator-endorsed policies for usage of data and sharing these policies with the information custodians within the Alliance Federation; and facilitate effective, dynamic data security risk management, monitoring and auditing throughout the information lifecycle.

The SDS enables users (e.g., originators, producers, owners and custodians) with:

1. The automated labelling of aggregated data elements based on policy;
2. The redaction of data elements from the content release to each recipient based on policy;
3. The controlled routing of messages to specified communication channels based on policy;
4. The controlled administration of policies by authorized administrators;
5. The logging of all transactions processed by the SDS to enable monitoring and auditing; and
6. The tracing of operational policies (rules and constraints) to administrative changes, certified policy sets and/or mission architectures;

3.2.3.2 Protect

Provide universal (i.e. anytime, anywhere, any data) protection of data at rest, in transit and in use within a CIS by complementing the existing security protection mechanisms, such as emission security, network security, and traffic flow confidentiality, through applying a set of security measures at a granular level that is dynamically adaptive to changing operational requirements. The goal is to enable the Alliance Federation to achieve a required level of protection of information that is stored, processed or transmitted in the Alliance CIS with respect to confidentiality, integrity, and availability.

The SDS enables several protections for data, including:

1. The automated labelling of data based on content sensitivity and policy;
2. The redaction of sensitive data content (at each aggregation point) based on user needs, authorizations and policy;
3. The encryption of any part of the data set based on policy;
4. The traditional access and release controls; and
5. The clamping of individual data environments using secure operating systems, software firewalls and their policies.

3.2.3.3 Share

Provide the ability to manage information with an emphasis on the 'responsibility-to-share' (balanced by the security principle of 'need-to-know'), and to facilitate discoverability, access, use/re-use, and reduced duplication, all in accordance with security, legal and privacy obligations with respect to all data and all users within the Alliance Federation. Therefore, DCS facilitates secure, interoperable and timely access to relevant information for those that require it.

The SDS enables the user to define, implement and automate the enforcement of policies that govern the sharing of data to the attribute (schema leaf-node) level.

3.3 DESCRIPTION OF A SECURE DATA SERVICE (SDS)

As illustrated in Figure 4, the SDS aligns well with the DCS defence in depth concepts. The SDS provides a DCS solution for deploying Data-as-a-Service (DaaS), placing a security boundary around the specific data content.

The following items identify how the SDS delivers DCS to a data environment:

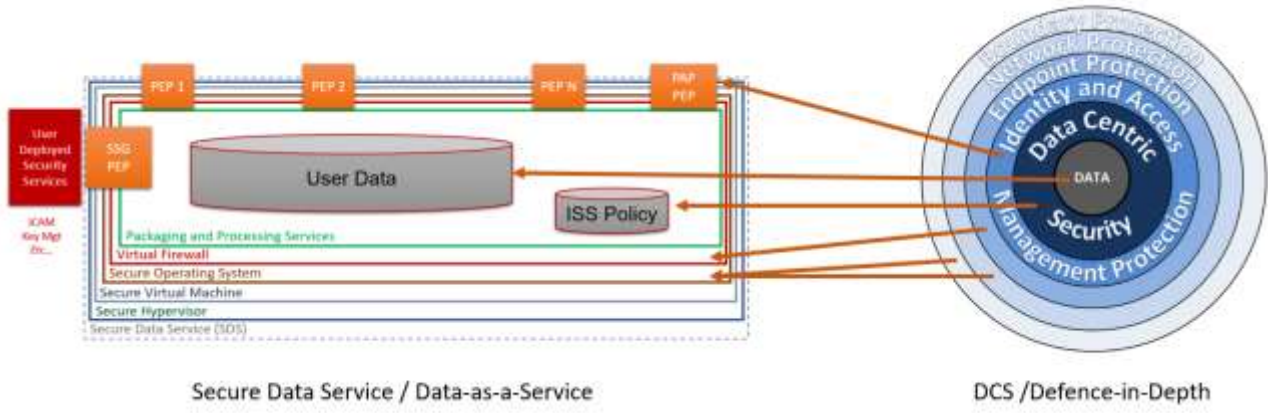

Figure 4:DCS - Alignment with SDS Concepts

1. **Data**: The SDS wraps the data to be protected and its storage services with a set of DCS services that establish defence in depth for the SDS protected data;
2. **Data Centric Security**: The SDS provides a set of packaging and processing services that arbitrate and enforce User (/Nation) defined policies (rules and constrains):
    a. **Processing**: Services that take information elements and parse, transform and marshal data elements to the data store;
    b. **Package**: Services that aggregate, transform, label, and redact information data and metadata elements in accordance with each recipients' information needs and authorizations; and

c. **Route**: Upon the completion of packaging, the metadata and data elements are formatted, embedded in the appropriate protocols and bindings, and routed to the agreed (/approved) exchange technologies and communication channels;

3. **Identity and Access Management Protection**: The SDS PEPs can integrate with the user specified Identity, Credential and Access Management services to control access and release of information;
4. **Endpoint Protection**: The SDS provides its own operating environment (Secure OS) where the user can install and operate their own selected endpoint security (antivirus, anti-spyware, application control and other styles of host intrusion prevention);
5. **Network Protection**: The SDS enables the user to define firewall policies to control network traffic and allowable communication protocols; and
6. **Boundary Protection**: As illustrated, the SDS places the boundary around the data store. This is important with the growing focus on cloud deployments, where boundaries are more difficult to identify and control.

3.3.1 SDS Deployment

The SDS enables users to tailor security features and policies to the sensitivity of the data being protected, rather than the network or operating environment as a whole. This permits users to deploy a SDS and associated user-data to any information domain, provided the ISS policy strictly enforces user (/national) and coalition sharing agreements, policies and doctrine.

As illustrated in Figure 4, this deployment can be achieved because the SDS is deployed with its own boundary, network, endpoint, access, and data protection features, and users can configure it to take advantage of other deployed capabilities. This is accomplished using:

1. The Policy Enforcement Points (PEP): PEPs provide the primary integration point between the SDS and its external environment. Users can implement a PEP to interface with multiple user systems and services or dedicate to a specific service. The PEP presents a common API to SDS services that can be used to integrate the SDS into the User or coalition environment. Features typically integrated into the PEP include: access controls, receipt controls, release controls and logging;

2. Security Services Gateway (SSG): Similar to the PEP in function, it provides a secure access to the user specified security services, including: Identity, Credential and Access Management (ICAM), Key Management (e.g., generation and escrow), and external Policy Decision Point (PDP). The SSG provides the SDS with dedicated interfaces to user specified and authorized security services;

3. Administration PEP: The Admin-PEP provides the SDS with dedicated interfaces to user specified and authorized management and administration services or Policy Administration Point (PAP). These services enable the specified users to configure the operating and policy environment of the SDS; and

4. Firewall: A software firewall is used to control access to ports and protocols to the SDS services. The firewall is included to minimize vectors of attack by only enabling ports and protocols used by the SDS Mission Configuration, and only to those systems, middleware and services authorized for the mission (including general day-to-day) operations.

For information on the functions, controls and interfaces for the PEP, SSG and PAP refer to the IEF-RA (Reference H).

3.3.2 Key Features

3.3.2.1 Separation of Concerns

The IEF is designed to separate concerns and increase the flexibility for key stakeholder groups:

1. Operator Concerns: Providing the ability to rapidly deploy and adapt capabilities to evolving operational needs, by enabling:
    a. Standard APIs:
    b. Reconfigurable services:
    c. Policy administration: and
    d. Policy Libraries:

    The SDS provides the ability to configure the service to operate (receive, process, store, package and release) on most structured and/or semi-structured data (native or semantic reference model). SDS policy environments can be developed for any data environment, and tailored to the ISS requirements for a specific mission and set of partners;
2. Business/Operational Concerns: Providing the ability to develop and deploy capability without having to define every business and technical requirement;
3. By separating the policy and software lifecycles, business, operational, information and security analysts can design, test and certify policies, and deploy them to operations without the cost and risk of a software development process. In addition, business can be more comfortable developing ISS capability to environments when they know they have the ability to adjust the policies.

    In addition, SDS services can be developed or enhanced without having to define every information exchange before a contract can be awarded – mitigating a significant risk for many major and minor capital projects;
4. Information and Data Management Concerns: Providing policy models using a standardized UML profile, IEPPV (Reference I), and aligning these models to other relevant architecture views and viewpoints (e.g., Strategic/Capability, Business/Operational, Information, Solution/System, Services, Data, Security and Technology); and
5. IT Concerns: Providing a reference architecture and services that can be securely deployed to on-prem, deployed platform, cloud and hybrid environments will increase Chief Technical Officers flexibility, agility and adaptability.

3.3.2.2 ISS Policy

The development of the Information Exchange Packaging Policy Vocabulary (IEPPV^{TM}, Reference I), truly enables the SDS. The IEPPV enables the user (e.g., data originator, owner, producer or custodian) to independently evolve:

1. The ISS policies (rules and constraints) independent of the software that enforces them;
2. The semantic patterns independent of the patterns governing its protection;
3. The semantic patterns independent of the patterns for their exchange; and
4. The data transforms independent of the semantic patterns.

This independence in development enables larger teams to use and reuse the policy patterns for variations in mission, Community of Interest (CoI), domain, system, application and individual needs and authorizations. This in turn mitigates key elements of risk and cost in interface design.

3.3.2.3 Policy Aligned to Architecture

The IEPPV provides a direct link and alignment to the Architecture Framework as it defines a profile for the Unified Modelling Language (UML™), which is one of the most commonly used modelling languages for software systems. The IEPPV also has an alignment with the operational interactions in the Unified Architecture Framework (UAF™) and the OV-2 and NOV in DODAF and MODAF respectively. This enables the DCS rules and constraint to be aligned with application interfaces, nodes, participants, systems, operations, missions, capabilities and strategy and more. This alignment provides much greater traceability and understanding of how data and information elements are used within the users' environment.

Enabling the definition of ISS policy in architecture models enables analysts (e.g., data, information, operations and security) to review policy without code reviews. It also allows the user to exploit Model Driven Architecture (MDA) and Model Based System Engineering (MBSE) to auto generate the runtime versions of the policies as an exchangeable data set, eliminating the long, costly and risky interface development cycles.

The data or metadata underpinning most architecture tools provides the opportunity to development analytic tools that users (e.g., management, security analysts and QA analysts) can employ to support multiple governance and audit activities (see Figure 5) starting during the design phase, and produce much of the documentation needed to develop and certify the capability.

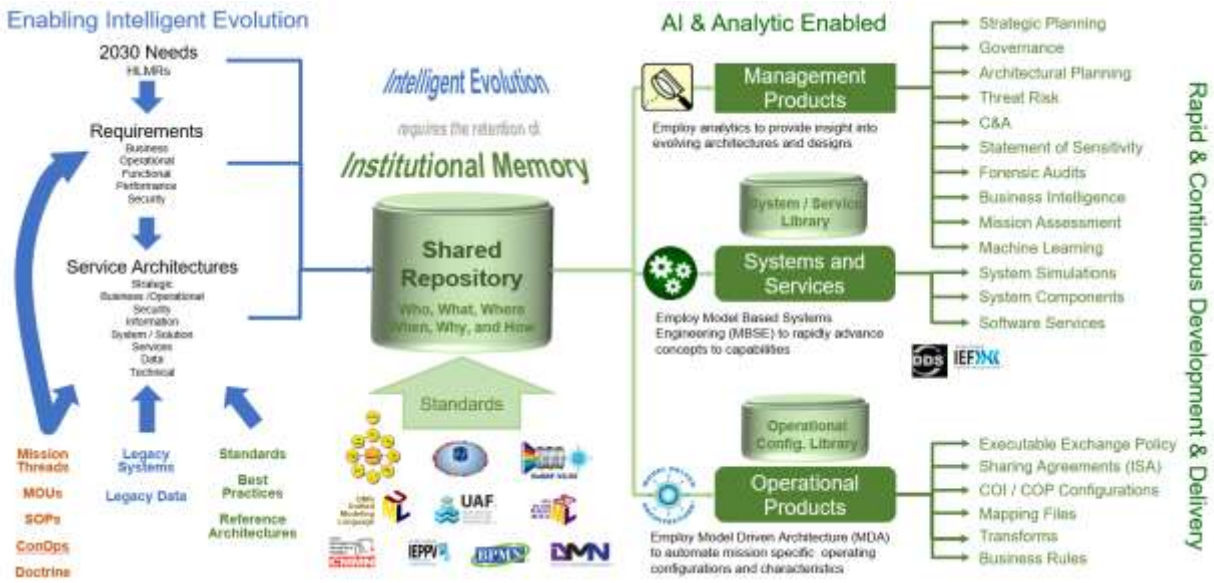

Figure 5: Analytics on Architecture for Governance and Audit

Finally, developing ISS policies within an architecture construct enables the retention of institutional memory about how data and information is used within the enterprise and in conjunction with mission partners.

3.3.2.4 Policy as Data

As illustrated in Figure 6, the Packaging and Processing Service is designed to import a series of ISS policy files (e.g., XML, JSON, or Binary objects), including (grey elements):

1. Semantic Policy;
2. Information Exchange Policy;
3. Message Schemas;
4. Data mapping files;
5. Decision logic; and
6. Configuration files.

Each of these files is managed and maintained by the user and reflects the ISS requirements for the specific mission, domain, or system. Updating or replacing the contents of the files, through the Policy Administration Point (PAP) interface (/PEP) enables the user to reconfigure the operation of the SDS to address changes in a mission environment, or configuration of the service for an alternate data environment or mission.

3.3.2.5 Reusable / Configurable Software Services

The SDS aligns and integrates several software services that can be structured into individual SDS configurations, which include:

1. The Semantic Processor (mandatory);
2. Information Exchange Controller (mandatory);
3. User and Generic Parsers (1 mandatory);
4. User and Generic Publishers (1 mandatory);
5. User and Generic Data Transformations (optional);
6. PAP PEP (Optional – if administration is authorized);
7. Security Services Gateway (optional – if integration with user cyber and security services is needed);
8. PEP (1 mandatory) – provides integration with the users' middleware;
9. CTS (Optional) – provides an integration with user specified cryptographic services; and
10. TLS (Optional) - provides integration with user specified logging services.

Not illustrated in Figure 6 are the Policy Decision Point (PDP) and Access Control Policies typically associated with a Policy Enforcement Point. For the purposes of this operating concept, these elements are considered part of the Users Access Management Environment. The SDS PEPs provide the integration point through which the SDS is aligned and integrated with the users' environment.

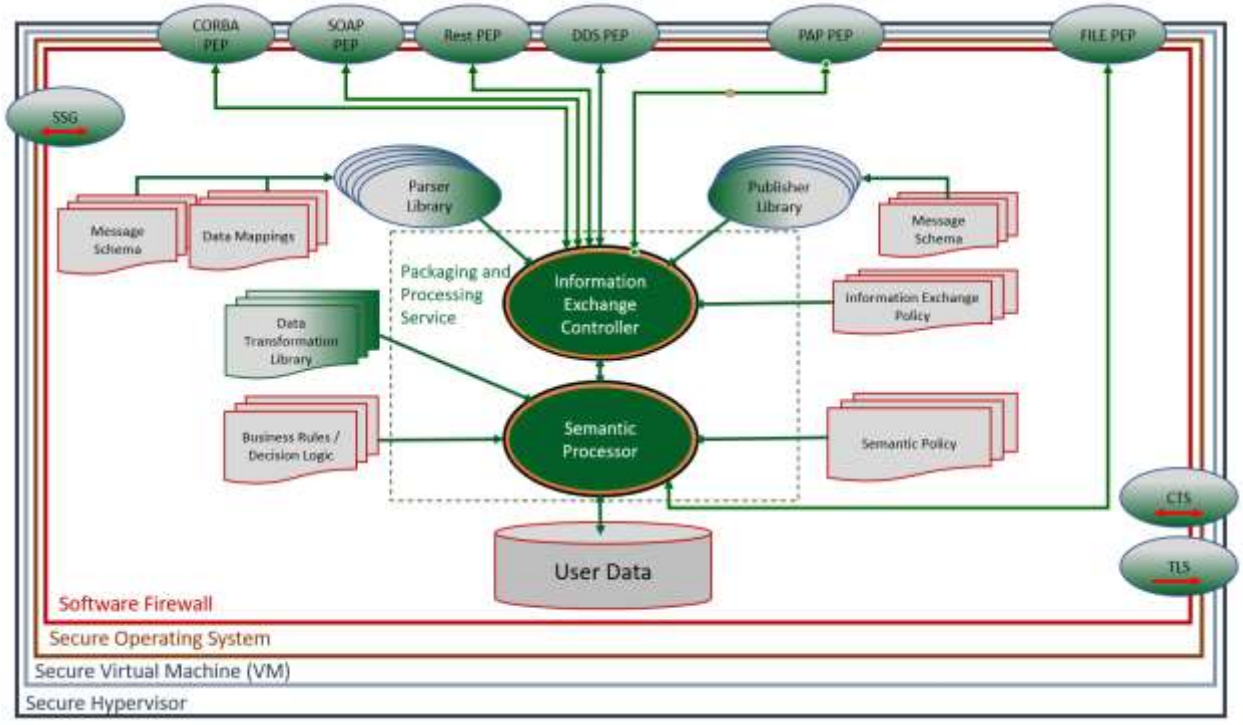

Figure 6: SDS Components

4. USERS OR INVOLVED PERSONNEL

4.1 OPERATOR

The typical operator will witness little or no direct interaction with the SDS. The SDS forms a layer between the operator applications (or user interface) and the data layer. Based in the implemented policies, the operator will be provided with the information needed and authorized to access using the interface they are accustomed to.

In some instances, operators may be authorized to perform administrative tasks related to the configuration of the SDS services and/or ISS policies. In this case, their user applications will be fitted with tools that interact with a PAP PEP configured with their authorized actions.

4.2 ADMINISTRATOR

The SDS provides a wide range of configuration and policy options, including:

1. Selection of data storage models;
2. Selection of storage technologies;
3. Selection of exchange models;
4. Selection of exchange technologies;
5. Selection of service configurations; and
6. Real time access to administration of the contents of the elements above providing the ability to adapt the SDS to changes in the mission environment.

The elements shaded in "grey" in Figure 6, if enabled, provide administrative control over their contents and therefore the features and functions of the services they govern, e.g.:

1. PEPs may be added or removed from a deployment configuration based on mission needs. During operations PEPs may be activated or deactivated by operators or administrators. The Decision Policies may be modified in accordance with the capabilities defined by the user specified Policy Decision Point (PDP). As with ICAM services the SDS is designed to integrate into the users' own security environment;
2. SSG is currently envisioned as a fixed configuration after start-up to mitigate the risk of the SDS being redirected to unauthorized security services;
3. CTS provides a link to the user specified cryptographic services or equipment; it is not using the software services within the SDS. As with the SSG the CTS is currently envisioned as a fixed configuration after start-up to mitigate the risk of the SDS being redirected to unauthorized security services;
4. The TLS will typically be configured at start-up;
5. Message Schemas and Data Mapping (MSDM) libraries hold the descriptors that govern how the SDS processes and marshals data received from external sources. Additional schemas and mappings may be loaded during operations provided they are consistent with the user data environment, transformation, and business rules and decisions logic present in the instance of the SDS;
6. Data Transformation Libraries (DTL) may be customized for a specific data store and augmented during operations. However, it is more likely the DTL will hold the entire set of transforms

employed by the user, and new or modified transforms will be added during development cycles for the SDS and data domain;

7. Information Sharing and Safeguarding Policies consists of:
    a. Information Exchange Policies: Derived from the IEF Information Exchange Specifications (IES) that are closely related to an electronic Information Sharing Agreement (eISA) in the IEPPV Specification (Reference I), but tailored to the operation of the SDS. The IES determines which messages can be received and processed by the SDS, and which data can be released formatted and released by the SDS;
    b. Semantic Policies: Also derived from the IEPPV (Reference I) and govern the aggregation, transformation, labelling and redaction of data and information elements for each recipient; and
    c. Business Rules and Decision Logic: Extends the IEPPV and increases the flexibility of data transformations and labelling operations, placing the rules under the control of the user rather than hard coded functions;
8. User Data is specified in the wrapper elements in the IEPPV (Reference H) and provides the ability for the user to specify how data elements are governed by technology (e.g., RDBMS, OODB, File Store, Object Store) selected.

For information on the functions, controls and interfaces for the PEP, CTS, TLS, and SSG refer to the IEF-RA (Reference H). Controls over User specified systems, services and equipment are not outlined in this document as they are subject to external user specification.

4.2.1 Analysts

In many instances, SDS policies and libraries are prepared by:

a. Operational Analyst: Defines the mission threads (e.g., use cases) and the information requirements of each commander (e.g., decision maker, operator, partner) in the thread;
b. Data Information Analyst: Defines the content, structure and source of data and information elements required by each user in a mission thread;
c. Data scientist: Defines the content and structure of the data elements needed by analytics tools and decision aids used by the commander; and
d. Security Analyst: Defines the restrictions on the release of data and information elements to each commander.

These elements are captured in the architecture and policy model for each mission and deployed SDS.

4.3 SUPPORT CONCEPT

The SDS is designed as a set of independent services, each with a defined interface or application program interface (API). The objective is to provide for a continual development environment where each service can be maintained or replaced over time by a small development team of Subject Matter Experts (SME), developers and quality assurance personnel. They would engage in agile development employing Development, Security and Operations (DevSecOps) to rapidly design, develop and deliver new services and capability to operations.

Separating the policy life-cycles from the SDS services that adjudicate and enforce them enables users to develop and test (e.g., desktop exercises) ISS policies and then deploy them to operations as and when needed. Implementing a library of ISS policies based on mission profiles, mission threads, and internal and partner

roles and responsibilities will enable increasing levels of reuse and the opportunity for Day-0 capabilities. This approach also enables the recapture of mission ISS deployments and mission-imposed changes as new or evolved ISS capability, further evolving internal and coalition interoperability.

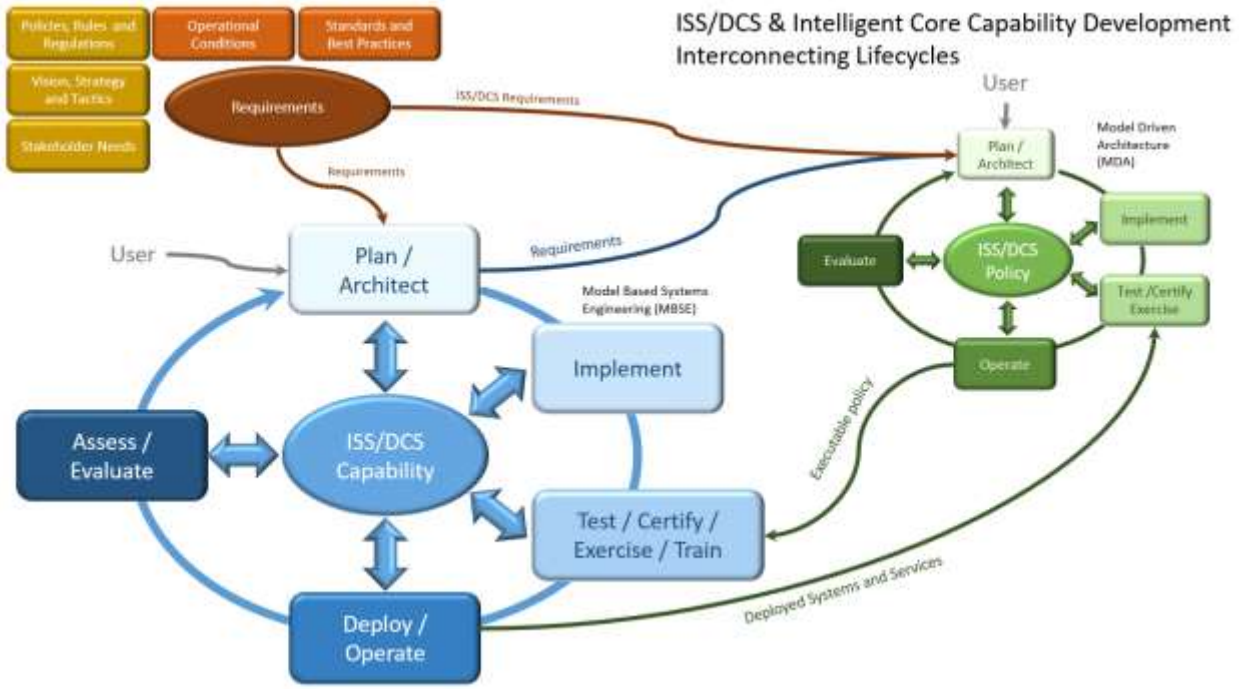

Figure 7: SDS Lifecycles

5. JUSTIFICATION FOR AND NATURE OF CHANGES

5.1 JUSTIFICATION FOR CHANGE

The primary reasons for developing this solution include:

1. The growing and evolving desire of organizations to all-source exploit data to better inform decision-making processes, and create a decision advantage;
2. The growing and evolving requirements for data and information requires services that can be rapidly configured and securely deployed to any environment (on-prem, deployed platform, cloud and hybrid environment) assets where they are needed – Data-as-a-Service;
3. The continual changes in a typical operational environment demands data services that are flexible, agile and adaptive. Services that can be configured in theatre by an authorized operator, or through commands and data from an authoritative source (e.g., higher headquarters);
4. Information sharing and safeguarding transcends any single system or software development project; and
5. The inability of users to articulate all ISS requirement and write the specifications needed to initiate a traditional development project.

Current information systems have developed, emerged, or evolved with varying degrees of independence, largely based on local needs and/or specific operational requirements. When designed and implemented they did not, or could not (e.g., out of scope requirements), take account of the broader enterprise requirements for operational integration and information interoperability. This has resulted in these solutions operating as stovepipes, providing only limited levels of information sharing and safeguarding capability. In addition, the interfaces to these systems are designed for peer-to-peer communication in a manner that is rigid and brittle and unable to adapt in alignment with operational tempo, and/or changes in the mission environment, e.g., changing threats, role, responsibilities, information needs, and partners. The SDS is specifically designed to:

1. Enable users to evolve ISS capabilities as needs and requirements are discovered;
2. Enable change throughout the ISS lifecycle;
3. Enable the secure and rapid deployment of ISS capability; and
4. Enable Data-as-a-Service.

5.2 DESCRIPTION OF NEEDED CHANGES

The SDS focusses on delivering data-as-a-service. This might require changes to existing information systems, refocusing users on a service-oriented solution.

To make full use of the increased precision that the SDS provides in the adjudication and enforcement of ISS requirements will require improvements to traditional information and data management practices. This too can be evolved over time and policy models can be developed and tested incrementally.

5.3 PRIORITIES AMONG THE CHANGES

The implementation of a core set of services:

1. Semantic Processer;
2. Information Exchange Controller;
3. One PEP integration to the users' middleware of choice;
4. The PAP PEP to enable administration of the SDS during operations;
5. Development of a policy environment for a priority data domain;
6. Integration of the PPS, Firewall, Secure OS and user Data Store into a VM for deployment; and
7. Practices and procedures for operating the SDS.

All other capabilities can be evolved from that baseline configuration.

5.4 CHANGES CONSIDERED BUT NOT INCLUDED

As identified throughout this document, the total set of ISS requirements for any user based on mission, threat, partner configuration, roles and responsibilities, ISAs, and other factors is unknowable. The approach takes into account the need to rapidly develop services and policy environments as institutional knowledge of data and information needs evolve. This operating concept does not address the practices and procedures needed for:

1. Agile Development;
2. Policy Development;
3. DevSecOps (Development → Security → Operations) for both services; and
4. Standard Operating procedures for User environments.

These elements should be specialized and integrated into the users own practices, procedures, standards and tools.

5.5 ASSUMPTIONS AND CONSTRAINTS

The SDS assumptions and constraints driving SDS as a solution, include:

1. Traditional system development practices are too rigid, brittle and protracted to address the rapid changes in information needs across the enterprise;
2. As data is the "ASSET" underpinning all IM/IT solutions its use, sharing, and protection must be maximized throughout its lifecycle; and
3. No ISS plan or design will survive first contact with operations, and will need to be flexible, agile and adaptive to be of functional use.

6. OPERATIONAL USE CASES

The SDS use cases relate to its possible deployments within operations.

6.1 USE CASE 1: USE OF SDS AS PART OF AN APPLICATION

The SDS in its most basic configuration (Figure 8) can operate as the data layer of a traditional 3-Tier of N-Tier information system architecture. In this case the ISS policies would be specifically tailored to the user services and the specific authorizations of the applications or system users. The Selected PEP would interface directly to the applications based on the interfacing technology chosen by the user, developer, vendor or integrator.

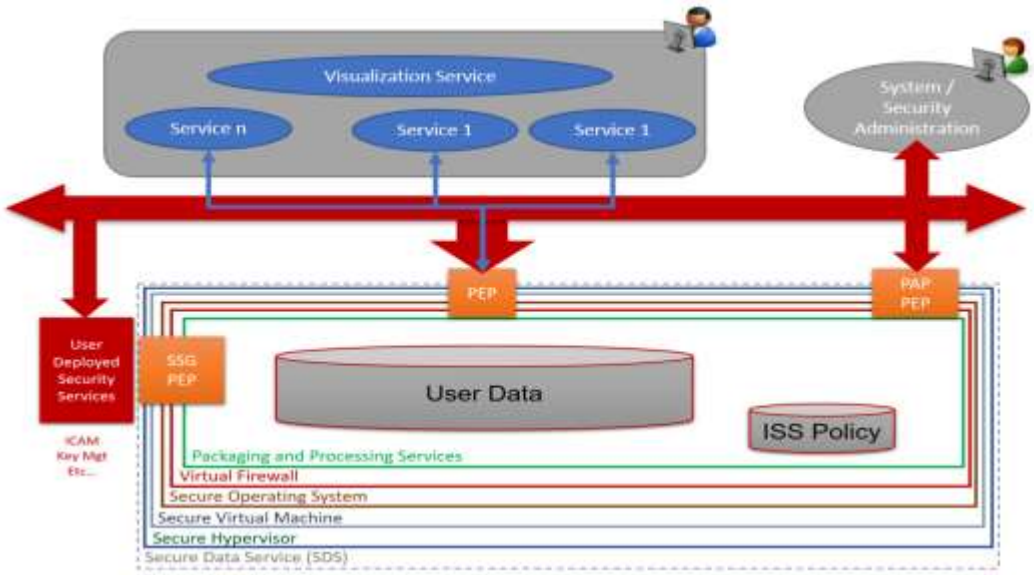

Figure 8: Traditional 3-Tier Architecture

6.2 USE CASE 2: USE OF SDS TO PROVIDE SECURE DATA-AS-A-SERVICE (sDaaS)

The SDS can also be configured (Figure 9) to operate as a data provider (e.g., Data-as-a-Service) for multiple concurrent data users (e.g., individuals, applications, services (e.g., data services), systems, organizations, and/or communities of interest). In this case, the SDS configuration may activate multiple technology interfaces (i.e., n*PEPs, SSG, TLS, and CTS) to meet the needs of a broad community of users. For this to occur, the user defines the configuration and policy environment required to address user information needs within the specific mission context.

The SDS services are being designed to be plug and play, reusable in and of the use case configurations, with their operation governed by user defined configurations and policy.

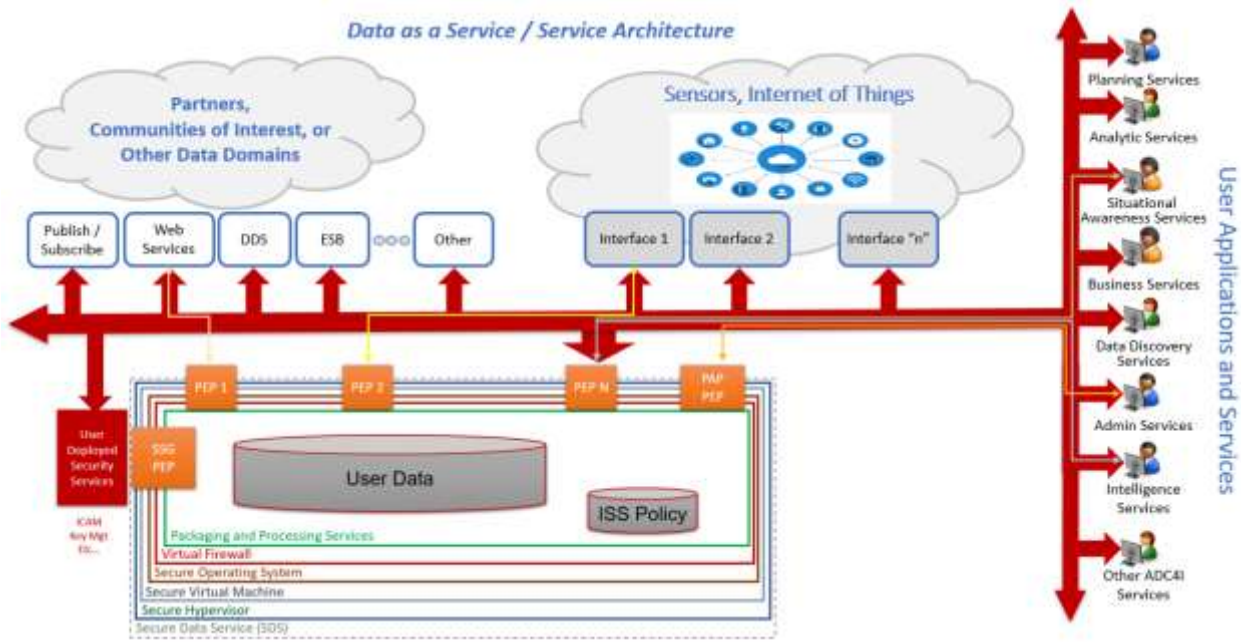

Figure 9: Data as a Service with DCS

6.3 USECASE 3: USE OF SDS TO DELIVER SELF SYNCHRONIZING DATA POOLS

In this configuration (Figure 10) multiple SDSs are combined to share and safeguard multiple data domains from those native to a specific sensor or application, translated into a normative form (e.g., Semantic Reference Model (SRM)) to enable down-stream analytics and decision aids, or as metadata to enable discovery and

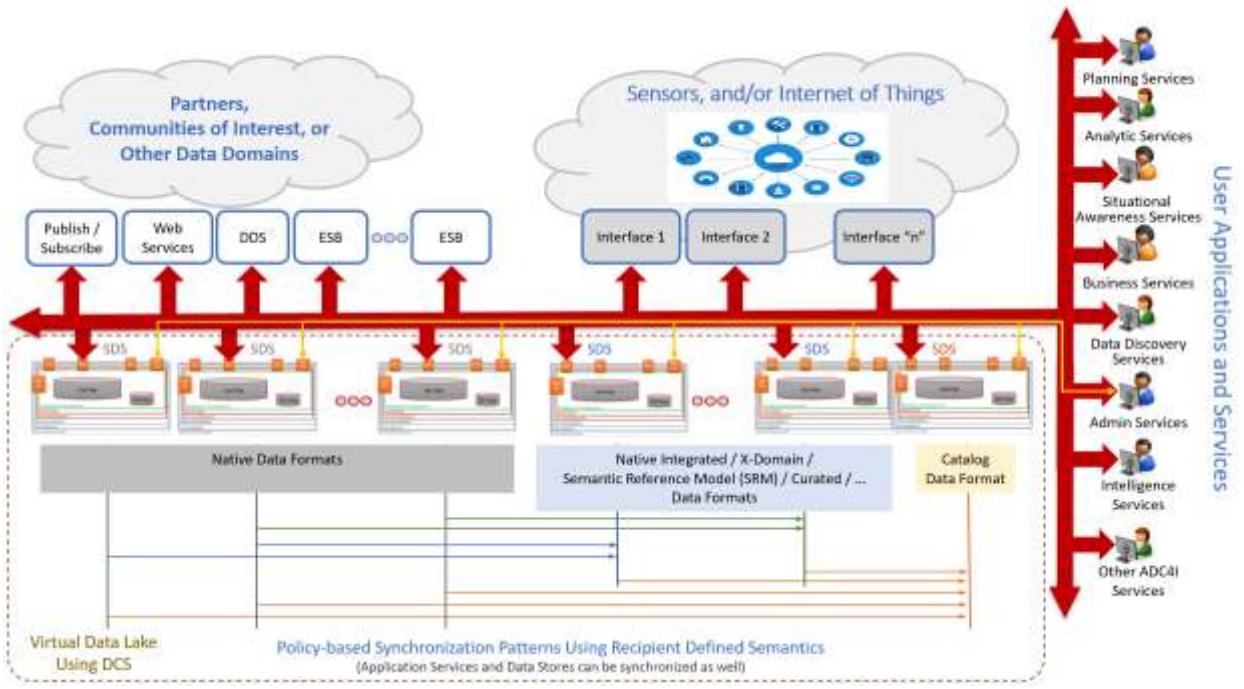

Figure 10: Secure Self Synchronizing Data Pool

access. In all cases, there is a single code base to support the data and technologies selected by the user or community to exploit available data assets.

In this way, the SDS becomes the foundational data service for a secure data and messaging fabric in an architecture that provides consistent capabilities across endpoints spanning on-premises, multiple cloud and hybrid environments. The SDS can consistently adjudicate and enforce user defined and controlled configurations and policies using a common set of software services.

The SDS can, through the implementation and deployment of PEPs, integrate and bridge:

- Multiples exchange technologies (e.g., REST, SOAP, DDS, ESB, and CORBA);
- Multiple data domains through the application of semantic (processing and packaging) policies tailored to the domain;
- Multiple users' needs and authorization through the overlay of filters and transforms in the semantic and information exchange policies; and
- Multiple data environments (differing exchange and storage semantics) by triggering sharing and safeguarding based on data event (data change or modification) triggers in the policies and variations in processing and packaging policy.

6.4 ALTERNATE SDS CONFIGURATIONS

This document outlines the scalability of the SDS (Figure 11) in terms of the use of multiple SDS instances within the data environment. Not discussed is the use of the elasticity of the cloud to add shared memory, processors and threaded versions of the SDS internal services (e.g., IES Controllers, and Semantic Processors) to scale processing capacity.

These configurations are planned for future development of the SDS core services.

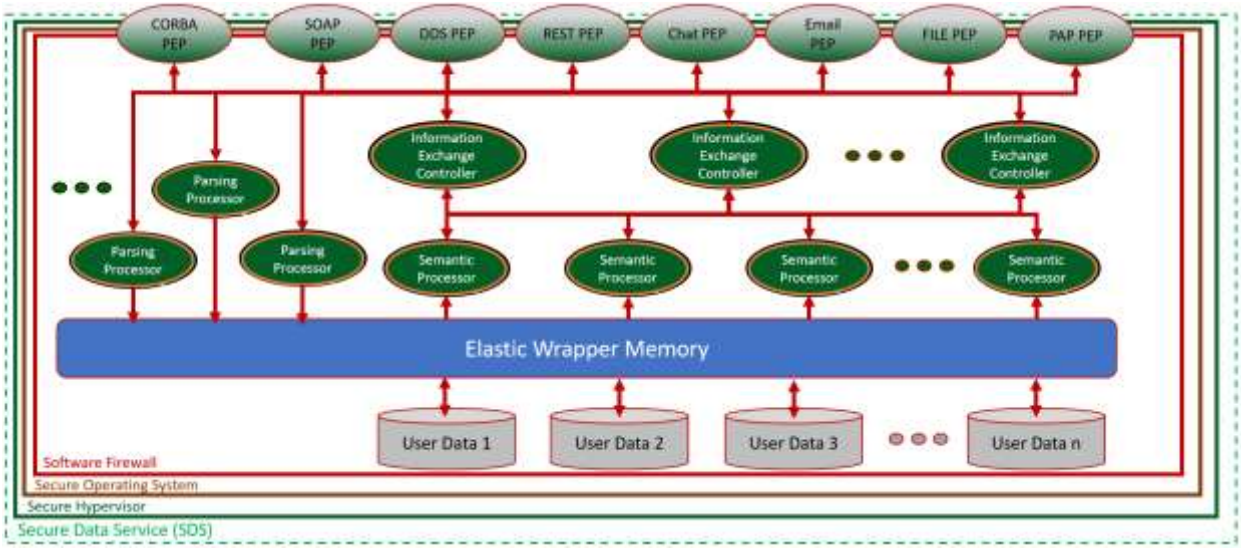

Figure 11: Elastic SDS

7. OPERATIONAL IMPACTS

7.1.1 Impact on Other Functional Services (Developers)

The objective is to develop and maintain the SDS using continual, agile, and service-oriented development practices to assure that it can continue to address changing mission and operational needs. New features and services will be developed, integrated and certified as users express and prioritise new requirements. Each development spiral would enable the SDS to employ new features, functions and technologies and the boundary points, e.g.:

- PEPs will be developed that integrate with additional middleware, Policy Decision Points and decision logic;
- SSGs will be developed to integrate with specific security services;
- The TLS will be developed to provide uses with logging technology options, e.g., Secure Database, SDS, and Distributed Objects (e.g., Block Chain) to secure the integrity of the logs; and
- Library elements, Schemas, business logic modules, data transformations, mapping files, parsers and publishers will be developed to interoperate with new and existing partners across new data domains as they evolve. These elements are considered library objects because they will be added to SDS configurations based on mission needs.

As the SDS matures, the number of changes to the core services (e.g., Semantic Processor, Information Exchange Controller and their interfaces to the PEPs, SSG, CTS and TLS) will diminish, changing only when new or extended policy types or features are added.

Development teams will be small and targeted to a very specific set of prioritized development requirements (e.g., a single PEP) that build on pre-existing services and interfaces (e.g., SDS internal messaging services based on the IEF Secure Message Bus).

7.1.2 Data Architects and Analysts

Post deployment, most of the development activity will focus on the development of executable ISS policy (e.g., IEPPV policy models) that governs how the SDS will transact on data and information elements. The approach is based on the IEF™, specifically the IEPPV (Reference I). Data analysts, architects and scientists will lead the development of ISS policy models that are aligned with enterprise and mission architecture views and viewpoints that demonstrate ISS alignment to the organizations, capabilities, mission threads, systems, applications, and interfaces. This alignment to architecture will enable users to assess and confirm that user data and information capture, processing, storage, analysis, visualization, sharing and safeguarding needs are met.

To enable and facilitate the data analysts, architects and scientists, the SDS is seeking to employ Model Driven Architecture (MDA) and Model Based Systems Engineering (MBSE) to translate the policy models from conceptual, to logical and on to operational capability. Where possible, the architecture tools will be extended to (auto) generate an operational (/executable) version of the ISS policies that can be tested in a lab environment, or directly releasable to operations. This automation will enable the user to eliminate teams of

API developers and the risks and costs associated with API development and maintenance. API responsibility (above technical integration) transfers to the SMEs and IM/DM groups within the users' enterprise.

This puts policy development into a separate life-cycle from the SDS software services. The separate life cycle policy development will be under the control of the operational community, providing this community with the management of that lifecycle, and assuring development is prioritised in accordance with mission needs.

Enabling and facilitating the operational community will require the implementation of new practices, procedures, standards and tools for the management of this policy-driven data-centric environment.

7.1.3 Impact on Other Functional Services (Users)

The SDS provides a high level of flexibility, agility and adaptability, many things configurable by the operator and/or administrator. The added flexibility, agility and adaptability will change how users think about the continual development and deployment of capability. They will need operating procedures, user interfaces and tools to manage and administer SDS configurations and operations in theatre. Functional Users (operators) will have to transition from the traditional project (fire and forget) requirements definition to full engagement in the definition, prioritization and testing of evolving ISS capability.

7.1.4 Impact on Support Agencies

The SDS is a highly configurable capability that will require training of support personnel to manage SDS configurations for deployment. They will also require:

1. The implementation of service configurations and policy catalogues that support personnel to configure a SDS for a mission and data domain;
2. The implementation of procedures and tools to recover architecture, configurations, policy and other artifacts to be catalogued and stored as new or enhanced capabilities (e.g., Day-0 capability);
3. The implementation of procedures and tools to rapidly augment, enhance or correct an ISS deficiency in service, configuration and/or policy identified during operations; and
4. The implementation of procedures and tools to securely and rapidly deploy new or enhanced capabilities to operations.

7.1.5 Impact on Operational Decision Making

Decision makers need to better articulate the types of information they need to enable and expedite the decision-making process and mission threads. These needs are then provided to the information and data architects, analysts and scientists so they can refine and specify the flow of data and information elements from source to the decision-maker, and what processing (e.g., collection, curation, analytics and presentation) needs to be exercised to provide the user with the information they require.

7.2 ORGANIZATIONAL IMPACTS

The implementation and deployment of the SDS will require extended skills and competencies in the Information Management and Data Management organizations, including:

1. ISS Architecture;
2. ISS policy development, testing and certification;
3. Information and Data Architecture;
4. Data protection rules and constraints and their alignment to users/partners, roles, missions and architectures;
5. Data science related to the collection, use, storage and analysis of data; and
6. Reviews and audits from design to operation.

7.3 IMPACTS DURING DEVELOPMENT

The SDS is being implemented as a set of continual standards-based services and API development. The objective is to develop the SDS and supporting ISS capabilities using the following continuous development practices (or equivalent):

1. Scaled Agile Management where strategic, operational and tactical requirements can be added to backlogs and prioritized for the development teams;

Agile Development where reallocated requirements can be broken down into development streams (e.g., sprints) and small teams can develop and test the new functionality;

1. Desktop exercising of new features can be assessed for viability and use of the services during missions; and
2. DevSecOps providing rapid deployment of approved and certified services to operations.

The development teams will need to develop skills and competencies in DCS and continual development practices, standards and tools.

7.3.1 Impact during Specification Development

As indicated, the SDS works best when on a continual development cycle, where small sets of requirements are allocated to individual services or micro-services. SDS capability is generated by configuring the core components (e.g., Semantic Processor and Information Exchange Controller) with a set of these services (e.g., library elements, PEPs, SSG, CTS, and TLS) for a specific mission or function.

In this environment, specifications include a small set of defined and constructed (testable) requirements that can be handed to a small development team to implement and test in a short amount of time. Backlogs of requirements are refined and priorities assigned by users (operators or SMEs) in conjunction with the development team. Users will be required to test, evaluate and accept the results of each development cycle. This too, will require some new skills and competencies to be developed in the user community.

7.3.2 Impact during Implementation

The benefits of the SDS are derived from the ability of components to evolve through their own lifecycles, governed by a set of principles and standards. Implementation relates to users configuring the SDS elements (e.g., Services, libraries and policies) for a specific mission deployment.

7.3.3 Impact during Operations

The SDS is an adaptable capability that lends itself to continuous development and DevSecOps. It will enable authorized administrators to rapidly adapt operational ISS to changing mission needs. If necessary, to request ISS policies or service components from a higher organization (e.g., Headquarters), national agency or approved libraries to augment, enhance or enable ISS within the mission.

Definition through operations is a continual process for:

1. Service features:
2. Supporting Elements:
    a. Policy Enforcement Points that integrate new middleware and communications capabilities;
    b. Data Transformation Library elements;
    c. Data parsers and mapping files;
    d. Data publishers; and
3. ISS policies:
    a. Information Exchange Specifications; and
    b. Semantic Policies:
        i. Processing;
        ii. Packaging; and
        iii. Business Rules.

Governance for deploying new features to operations must be developed and administration tools for these features need to be developed and deployed. Each feature can follow its own DevSecOps lifecycle (definition, implementation, testing, certification and deployment), independent of other components, but not beyond the standards that bind the SDS components.

8. ANALYSIS OF THE PROPOSED SYSTEM

8.1 SUMMARY OF ADVANTAGES

The SDS provides a single, reusable DCS solution that can be reconfigured for multiple data domains and deployments (on-prem, deployed platform, cloud, and hybrid). This will enable stakeholders:

1. An independent software service that can be deployed into multiple domains and environments, reducing training and maintenance risks and costs;
2. A set of software services that can be maintained within their own life-cycles reducing risk and cost;
3. Mission ISS policy libraries, parser libraries, publisher libraries, transformation libraries and other reusable components that can be configured, tailored and reconfigured to deliver a wide range of mission capabilities;
4. To adapt ISS operations to specific mission needs and constraints using certified library elements and service components;
5. Secure deployment to multiple environments; and
6. To evolve capability using small independent teams to evolve and broaden ISS capabilities employing projects that are not too big to fail or stop.

Based on open international standards, multiple implementations and integrations will become available from multiple vendors, suppliers, integrators and possibly open-sources – reducing risk and cost.

8.2 SUMMARY OF DISADVANTAGES/LIMITATIONS

The adoption of the SDS will require data producers and owners to improve their own understanding of their data environment and information environment, which involves practices such as:

- Information Management;
- Data Management;
- Information and Data Architecture;
- Information Sharing Agreement Management;
- Data Operations; and
- ISS policy development.

Current Information systems typically operate with dedicated data stores versus Data-as-a-Service. This may impact on the System development:

- COTS options may be restricted due to their inability to work with DaaS; and
- Development practices may need to be updated to employ DaaS and service-oriented solutions.

8.3 ALTERNATIVES AND TRADE-OFFS CONSIDERED

Other approaches considered are:

- Current State:
    - These solutions do not deliver the desired levels of interoperability;

  - These solutions do not deliver the desired levels of security; and
  - These solutions do not provide the desired levels of operational flexibility, agility and adaptability;
- System API Based DCS:
  - These solutions typically require a rigid set of policies to be embedded within its code base – making the approach too rigid and brittle for real-world operations;
  - Each change to the API requires an iteration through the SDLC that is also fairly rigid and brittle;
  - Part of application-based security – maintaining the data as part of specific applications will hinder the ability of stakeholders to exploit all domain information; and
  - Does not deliver the flexibility, agility and adaptability sought by stakeholders; and
- Boundary PEPs:
  - These solutions work only on metadata bound to the message and/or payload, and as indicated, few information systems effectively automate the labelling process and S2S operating at machine speeds makes manual labelling impractical;
  - These solutions typically offer a go-no go option for the release of information – there is no capacity to redact information or data elements to assure that recipients receive some information necessary to effectively and efficiently render a decision; and
  - Does not address the balancing of sharing and safeguarding sought by stakeholders.

Therefore, the SDS is selected as the approach most likely to succeed, and deliver on stakeholder objectives.

8.4 OTHER CONSIDERATIONS

8.4.1 Governance

Providing the SDS levels of flexibility, agility and adaptability will require new governance practices, processes, standards and tools to enable auditors to analyse and assess all phases of definition, design, implementation and operations. This will require tools ((see Figure 5) that automate the analysis of architecture, design, test and operations data to develop the artifacts traditionally required to govern information and data management operations, including:

- Statements of Sensitivity;
- Threat-Risk Analysis; and
- Security and Forensic Audits.

8.4.2 Auditing

Highly interoperable systems operating in multiple data and security domains will require higher levels of:

- Runtime/real-time monitoring;
- Alerting; and
- Forensic auditing.

8.4.3 Certification

Employing an agile development and DevSecOps to rapidly develop and deploy operational capabilities will require the ability to execute and approve delta certifications for new or enhanced services/micro-services. This will require new security policies, practices, processes and tools.

ANNEX A

The following definitions are used within the SDS OCD.

Agile Development	Practice approach discovering requirements and developing solutions through the collaborative effort of self-organizing and cross-functional teams and their customer(s)/end user(s). It advocates adaptive planning, evolutionary development, early delivery, and continual improvement, and it encourages flexible responses to change.
Application Program Interface	Definition of the rules, constraints and configuration governing interaction with the host application.
Data	Facts (such as measurements or statistics) used as a basis for reasoning, discussion, or calculation.
Data as a Service	Information provision and distribution model in which data is made available to consumers over a network environment.
Data-Centric	The adjudication and enforcement of information sharing and guarding policies (rules and constraints) governing individual data and information elements.
Data Lake	System or repository of structured, semi-structured or unstructured data stored in its natural or raw format using a flat architecture (a data warehouse is a repository for structured, filtered data that has been processed for a specific purpose).
Day-0 Capability	A set of services and/or resources that can be employed to address or mitigate an incident, event or vulnerability on the day of discovery.
DevOps	Practice that combines software development (Dev) and IT operations (Ops). It aims to shorten the system development life cycle and provide continuous delivery with high software quality.
DevSecOps	Integration of security evaluation and testing at every phase of the software lifecycle, from initial design through integration, testing, deployment, delivery and maintenance.
Forensic Auditing	Ability to analyse the architectures, designs and/or operational logs to verify that components are operating properly, and effectively enforcing information sharing and safeguarding policies appropriately.
Identity, Credential and Access Management	Service to control access and release of information based on individual user authorisation and need to know.

Information	(1) Data in context; and (2) Data in a form that informs a decision.
Information Exchange Specification	Exchange specification between two or more parties specifying how information is to be shared between each party (equivalent to the Information Sharing Agreement used by the US).
Information Exchange Framework Reference Architecture	An OMG sponsored open reference architecture for information sharing and safeguarding, employing data centric security principles.
Information Exchange Packaging Policy Vocabulary	Vocabulary that will provide consistent concepts for the expression of rules governing information packaging and processing.
Information Sharing Agreement	Exchange agreement between two or more parties specifying how information is to be shared between each party.
Information Sharing and Safeguarding (ISS)	A set of capabilities that provide users with the ability to responsibly share information based on user needs, user authorizations and data sensitivity.
Intelligence	(1) Understanding / comprehension of the available information; (2) Insight into the current situation; and (3) Assessment of future events or situations.
Memorandum of Understanding	Statement defining the specific criteria that forms the basis of the understanding between parties.
Model Based Systems Engineering	Systems engineering methodology that focuses on creating and exploiting domain models as the primary means of information exchange between engineers, rather than on document based information exchange.
Model Driven Architecture	Software design approach for the development of software systems providing a set of guidelines for the structuring of specifications which are expressed as models.
Operational Concept Document	Discussion paper describing the technical or operational need being addressed and the goals, objectives, features and functions of a proposed solution to address that need, along with an assessment of impact on user environment and operational use of the proposed solution.

Operational View-2	Applying the context of the operational capability to a community of anticipated users with the primary purpose of defining capability requirements within an operational context.
Packaging and Processing Service	Transition structured information elements between data stores and information exchange services in accordance with local information sharing and safeguarding policies.
Policy Administration Point	Provides an authorised user with an interface to access services needed to manage and administer the configuration and policy environments of IEF components.
Policy Decision Point	Adjudicates access to, or the release of resources to a specified user based on resource sensitivity, user privilege and operational context in which the decision is being made.
Policy Enforcement Point	An integration point between the User's infrastructure and the SDS service which enables the user to integrate access controls to the receipt and release of messages.
Policy Driven	The adjudication and enforcement of rules and constraints derived from, and traceable to, user or community approved policy instruments (e.g., legislation, international agreements, regulations, directives, information sharing agreements, operating policy and operating procedures).
Publish/Subscribe	Architectural design pattern that provides a framework for exchanging messages between publishers and subscribers. This pattern involves the publisher and subscriber relying on a message broker that relays messages from the publisher to the subscribers. The host (publisher) publishes messages to a channel that subscribers can then sign up to.
Request/Response	Message exchange pattern that generates a suitable response against a correctly prepared request.
Scaled Agile	A set of organization and workflow patterns intended to guide enterprises in scaling lean and agile practices to plan, prioritize and manage capability development. Scaled Agile enables an enterprise to expand Agile development practices beyond the application development process.
Security Services Gateway	Provides a secure access to the user specified security services.
Semantic Reference Model	A database model describing the structured entities found within the model and all the relationships that exist between them.
Software as a	

Service	Software licensing and delivery model in which software is licensed on a subscription basis and is centrally hosted.

ANNEX B

The following acronyms are used within the SDS OCD.

Acronym	Definition
AD-C4I	All Domain Consultation, Command, Control, Communications and Intelligence
AI	Artificial Intelligence
API	Application Program Interface
C2I	Command, Control and Intelligence
C4I	Command, Control, Communications, Computers and Intelligence
CFI	Connected Forces Initiative
CIS	Communication and Information System
CoI	Community of Interest
CORBA	Common Object Request Broker Architecture
CTS	Cryptographic Transformation Service
DaaS	Data as a Service
DataOps	Data Operations
Day-0	Day Zero
DCS	Data Centric Security
DDS	Data Distribution Service
Dev	Development
DevSecOps	Development, Security and Operations
DODAF	Department of Defense Architecture Framework
DTL	Data Transformation Library
eISA	Electronic Information Sharing Agreement
ESB	Enterprise Service Bus
FMN	Federated Mission Networking
HQ	Headquarters
ICAM	Identity, Credential and Access Management
IEF	Information Exchange Framework

IEF-RA	Information Exchange Framework Reference Architecture
IES	Information Exchange Specification
IEPPV	Information Exchange Packaging Policy Vocabulary
IM	Information Management
ISA	Information Sharing Agreement
ISS	Information Sharing and Safeguarding
IT	Information Technology
JSON	JavaScript Object Notation
MBSE	Model Based System Engineering
MDA	Model Driven Architecture
MODAF	Ministry of Defence Architecture Framework
MOU	Memorandum of Understanding
MSDM	Message Schema and Data Mapping
NAF	NATO Architecture Framework
NATO	North Atlantic Treaty Organisation
NCDF	NATO Core Data Framework
NGO	Non-Government Organisation
NOV	NATO Operational View
NNEC	NATO Network Enabled Capability
O&M	Operations and Maintenance
OCD	Operational Concept Document
OGD	Other Government Department
OMG	Object Management Group
OODB	Object Oriented DataBase
Ops	Operations
OS	Operating System
OV-2	Operational View 2 – Operational Resource Flow Description
PAP	Policy Administration Point
PDP	Policy Decision Point
PEP	Policy Enforcement Point

PPS	Packaging and Processing Service
QA	Quality Assurance
RDBMS	Relational DataBase Management System
REST	Representational State Transfer
S2S	System to System
SDS	Secure Data Service
Sec	Security
SDLC	Software Development Life Cycle
SME	Subject Matter Expert
SOAP	Simple Object Access Protocol
SOS	Secure Operating System
SRM	Semantic Reference Model
SSG	Security Services Gateway
STANAG	Standard NATO Agreement
STF	Standards Transformation Framework
TLS	Trusted Logging Service
UAF	Unified Architecture Framework
UML	Unified Modelling Language
VM	Virtual Machine
XML	Extensible Markup Language