accenture

consulting | technology | outsourcing

Blockchain Traceability for the Counterfeit Detection and Avoidance Program (CDAP) Final Report

High performance. Delivered.

Date: July 17th, 2020
Point of Contact: Randy Gowat (randall.c.gowat@accenturefederaldefense.com)

Table of Contents

List of Tables

List of Figures

Executive Summary

The Counterfeit Detection & Avoidance Program (CDAP) plays an important role in ensuring the Defense Logistics Agency (DLA) procures critical electronic components from reputable vendors and manufacturers. The program includes pre-qualification of vendors and a post-award review process to ensure vendor documentation and traceability satisfies risk mitigation thresholds (i.e. that the documentation is sufficient to conclude the components are safe to procure). However, this process is highly manual and requires a large amount of additional correspondence with vendors concerning administrative mistakes or omissions in their submissions.

Starting with the Digital Traceability (DT) project, the CDAP processes were analyzed and an enhanced digitized methodology was proposed that showed potential to provide the program with automation and increased efficiency. Further, a feasibility study was conducted to assess whether emerging blockchain technology could be utilized to introduce anti-fraud capabilities as a backend component. This feasibility study demonstrated the potential value blockchain could bring to the program, leading to the inception of this phase of the project – the Blockchain Traceability for CDAP Prototype.

Over the course of the past several months, a diverse stakeholder group was formed including component representatives across CDAP, programmatic leaders across several Major Subordinate Commands (MSCs), leaders from DLA Headquarters, third-party partners, and external vendors. This Trusted Working Group (TWG) put forth their vision for an improved and more efficient method of collaborating to complete CDAP requirements and these were captured as 'recommended requirements' for the prototype application.

Based on the set of collected requirements, a general design for the future-state CDAP solution was generated and a variety of blockchain development platforms (software packages) were analyzed and evaluated to determine which would provide the most effective solution for DLA. Ultimately, Hyperledger Indy was selected as the platform of choice and Hyperledger Aries and Hyperledger Fabric were noted for consideration and implementation where relevant.

Moving forward, a technical team was formed and began developing the prototype application with key tenets of effective software design in mind, and Agile frameworks including bi-weekly (every two weeks) demonstrations of in-progress results to the TWG. This iterative process allowed the development team to demonstrate progress and collect feedback from the stakeholder groups as the prototype moved from frontend concept designs to a finished end-to-end application.

The final prototype application includes an identity-focused blockchain that provides several novel features including (1) a near real-time credential verification button, (2) immutable records of vendor qualifications and related documentation, and (3) a process for onboarding vendors with a blockchain-based decentralized identifier (DID). These features, along with several other quality of life improvements such as automated email services, field-level validations, and help text provided the CDAP stakeholders with a greatly enhanced digital process compared to the current state. Additionally, enterprise considerations surrounding blockchain such as cybersecurity, governance approaches, and technical details of the cryptography involved were provided to inform DLA decision makers.

1.0 Project Background

This 9-month short-term project, ***Blockchain Traceability for the Counterfeit Detection & Avoidance Program (CDAP),*** was initiated to prototype a blockchain-based web application to continue previous work done with the CDAP Team. The CDAP processes and workflows were first analyzed as part of the ***Digital Traceability*** project and that work progressed with a blockchain-specific feasibility study. This feasibility study laid the groundwork for a DLA blockchain prototype and described the stakeholder processes and working groups in detail. Where appropriate, previous descriptions of the stakeholder groups, CDAP process, and related considerations will be employed within this document to provide as much relevant detail as possible. Additionally, this ***Final Report*** will aim to provide a comprehensive executive-level report of the project development and findings, with detailed supporting information delivered separately (or in Appendices).

1.1 Counterfeit Detection & Avoidance Program (CDAP)

Reports published by the Department of Defense and Government Accountability Office (GAO) make the threat of counterfeit and nonconforming parts entering the DoD supply chain abundantly clear. These reports identify electronic components as major targets for malicious activities due to their abundance in critical systems and their ability to be maliciously programmed or otherwise compromised. In response to this increased threat, DLA has tasked its CDAP team to conduct an additional analysis and review prior to the acquisition of the highest-risk electronic components. Currently, the CDAP process and additional protections are applicable for every acquisition of federal supply class (FSC) 5962 components (Electronic Microcircuits). Vendors who are awarded contracts in this FSC are required to meet certain levels of qualification and must provide material evidence of their qualification in the form of test reports or traceability documentation prior to shipping the materiel to DLA.

1.2 CDAP Stakeholder Groups

The CDAP processes involve many different internal and external stakeholder groups. Externally, these stakeholders include 3rd party contractors Applied DNA Sciences (ADNAS) and vendors who may operate as Original Equipment Manufacturers (OEMs)/Original Component Manufacturers (OCMs), distributors, or other resellers. Internally, the program is overseen by a core CDAP Team who perform post-award reviews and a variety of supporting departments including the Test Lab, Warehousing, and Procurement.

1.2.1 Core CDAP Team

The CDAP program is managed and operated by a team of J3 resources based out of DLA Land & Maritime (L&M). This team is responsible for communicating with vendors who have won awards within covered supply classes (only Federal Supply Class 5962 currently) to determine their qualification level and ultimately review and approve or reject supplemental documentation. This team is the primary Technical Quality (TQ) group responsible for communication with external vendors and coordinates with other teams throughout the CDAP process to track procurements throughout their CDAP lifecycle.

1.2.2 Warehousing Team

The Warehousing team, sometimes referred to as 'Defense Distribution Warren Ohio' (DDWO), is the group that directly receives shipments of FSC 5962 components from vendors. This group, which is co-located with the Test Laboratory Team onsite, receives shipments and performs the ‘quick look’ process. The ‘quick look’ consists of comparing package contents with cover sheets and procurement information to ensure the items match what is expected and are packaged properly (in most cases, this means enclosure in static-resistant packaging). Once the warehousing team has determined that a shipment passes ‘quick look’, they move the items on to the Test Laboratory to continue the CDAP process.

1.2.3 Product Test Laboratory Team

The Test Laboratory team is responsible for inspecting electronic microcircuits received through the CDAP process and for applying appropriate markings if the components pass inspection. The Test Laboratory receives items that have passed the ‘first look’ from DDWO and performs x-ray inspections on all received components. Once a microcircuit has passed the x-ray inspection, it is marked with deoxyribonucleic acid (DNA) provided through external partner Applied DNA Sciences and tracked within Test Laboratory databases.

1.2.4 Applied DNA Sciences

Applied DNA Sciences (ADNAS) is a third-party contractor DLA has partnered with to help execute portions of the ‘DNA marking process’ employed at DLA’s Product Test Laboratory. ADNAS produces unique proprietary DNA strands which are stored in ink and utilized to mark components received through the CDAP program. After a component has been received and gone through appropriate inspections, it is marked with specific ink and baked to ensure adherence. ADNAS also provides a quality control function by receiving and inspecting representative marking samples from each batch of items marked by the DLA Test Laboratory. Currently, ADNAS and the Test Lab share some level of marking information with each other for item tracking purposes in the case of a product recall or other need.

1.2.5 Original Equipment Manufacturers / Original Component Manufacturers

DLA prefers to obtain components received through CDAP from Original Equipment Manufacturers (OEMs) or Original Component Manufacturers (OCM) to ensure the quality and effective operation of these components. As a risk-mitigation program, it benefits CDAP to receive components directly from their manufacturing source as opposed to obtaining them from distributors or resellers who naturally would have obtained them from another source. For these reasons, vendors qualifying as OEMs/OCMs are among the most trusted DLA suppliers and as such require the least burden of proof when submitting items through CDAP with appropriate documentation. During this prototype phase, DLA partnered with Rochester Electronics to represent the role of an OEM/OCM within the requirements and testing phases of the project.

1.2.6 Distributors & Resellers

In an effort to meet the demand for microelectronics across the enterprise, DLA also periodically obtains components from distributors and resellers. These vendors are naturally a minimum of one degree removed from OEMs/OCMs and must provide proof of traceability back to the manufacturing source (i.e. OEM/OCM from whom the components were purchased) during post-award CDAP reviews. Although they are not inherently as trusted, CDAP provides a certification process for distributors (including ‘Authorized Distributors’) and resellers, who must then provide several pieces of documentation depending on the specifics of a given award. In general, this documentation may include traceability back to the OEM/OCM in the form of shipping documents or invoices, and other documentation such as test reports or the results of quality inspections. During this prototype phase, DLA partnered with Forward Components to represent the role of a distributor/reseller within the requirements and testing phases of the project.

1.3 CDAP Processes

As previously described, the CDAP processes rely on a variety of stakeholder groups and can vary depending on the qualification level of the vendor from whom the components are being obtained. This process remains largely manual in that most of the information exchanges with vendors are done via United States Postal Services (USPS) mail or by electronic mail (e-mail). Similarly, internal information exchanges among DLA groups is completed via e-mail, post-award requests (PARs), or by checking a shared database called the 'CDAP Catalogue'. The information exchanges required to process the reviews and inspections for a typical procurement of electronic components through CDAP is depicted below.

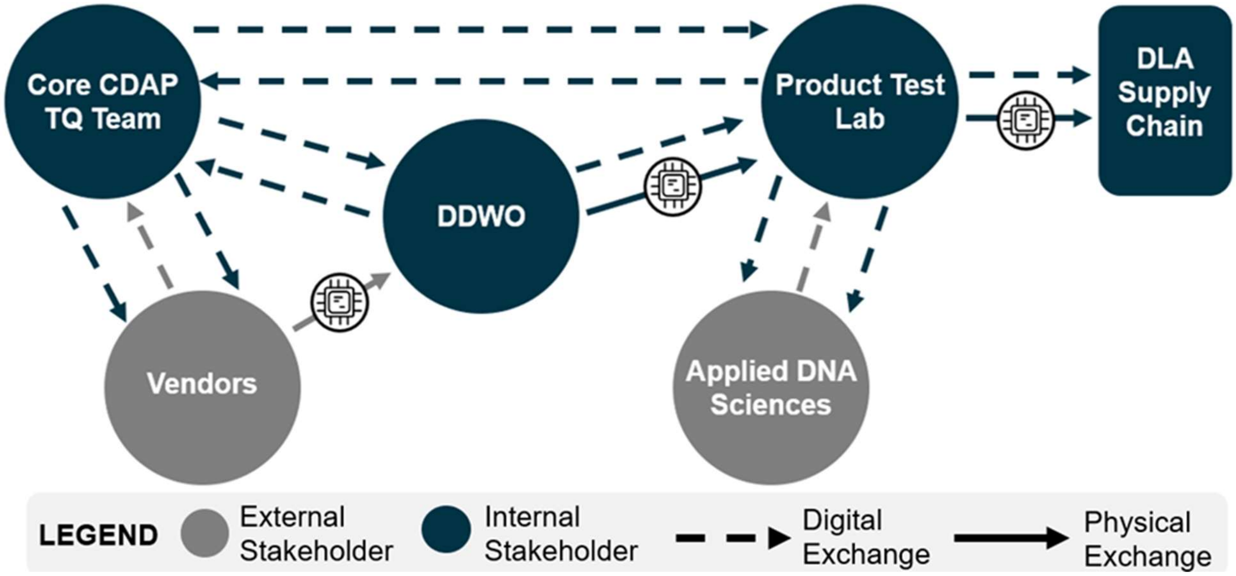

Figure 1. Typical CDAP Exchanges for One Procurement

The exchanges shown illustrate how complex and reliant the CDAP program can be on collaboration among stakeholder groups. However, it is also important to understand that vendors must first qualify for the program and attest to certain quality levels before they can even qualify for, and bid on, CDAP-related solicitations. This section will describe some of the typical CDAP processes so that the qualification levels and burden of proof required from vendors is made clear.

1.3.1 Pre-Award Qualification

Vendors must obtain a level of trusted qualification with DLA before being eligible to complete the requirements of an FSC 5962 award and deliver the materiel to DLA. It is unclear whether a vendor's qualification status is considered by contracting officers during the solicitation bid review process, but lack of any qualification will ultimately be a disqualifying factor if the awarded vendor cannot properly certify their merits with the CDAP team. There are five types of qualifications which would allow a vendor to participate in an FSC 5962 acquisition including:

(1) **Approved Source Manufacturer Specified in the Solicitation/Contract Item Description (Original Component Manufacturer (OCM), Original Equipment Manufacturer (OEM))** – the source identified by name, Commercial and Government Entity (CAGE) code, and part number listed in the purchase item description in Section B of the award. This entity only qualifies as OCM/OEM if they are the original manufacturer

of the part, the original component manufacturer, the original equipment manufacturer, the manufacturer specified in the solicitation/contract item description, or the part Design Control Activity (DCA), which is defined as the entity having responsibility for the design of a given part and for the preparation and currency of engineering drawings and other technical data (TD) for that part.

(2) **Approved Source on the Applicable Qualified Products List (QPL)/Qualified Manufacturers List (QML) –** requires a Certificate of Conformance and Traceability (CoC/T) and must include information and documentation required by the applicable military specification for the QPL/QML product, that is listed on the DLA Technical and Quality Requirement RQ007**.**

(3) **Authorized Distributor of the OCM/OEM or QPL/QML Approved Source** – A supplier who has a contractual arrangement or the express written authority with the OCM/OEM of the item being acquired to buy, stock, re-package, sell, and distribute the item specified in the solicitation/contract.

(4) **Supplier/Distributor on the Qualified Supplier List of Distributors (QSLD) for FSC 5961/5962** – A distributor listed or determined qualified for listing on the DLA Land and Maritime Qualified Supplier List of Distributors for FSC 5961/5962.

(5) **Supplier/Distributor on the Qualified Testing Suppliers List (QTSL) for FSC 5961/5962 –** A distributor listed or determined qualified for listing on the DLA Land and Maritime Qualified Testing Suppliers List (QTSL) for FSC 5961 and 5962.

The vendor types and descriptions come from information provided by the CDAP team as found on the current version of Form 918 (Traceability/Test Documentation Cover Sheet).

1.3.2 Award and Certification Qualification

Once a vendor has been awarded a contract to provide an FSC 5962 part, they must provide the CDAP office with proof, or certification, of their qualification in one of the five previously described qualification levels. This process must be completed at least fifteen days prior to the delivery date specified in the contract and requires a minimum of two documents be sent in for CDAP review: (1) a Traceability / Test Documentation Cover Sheet (also known as DLA Land and Maritime Form 918) and (2) Traceability proof or Test Report documentation. Information pertaining to the required documentation and a link to Form-918 are located on the DLA Land and Maritime CDAP webpage (http://www.dla.mil/LandandMaritime/Business/Selling/Counterfeit-Detection-Avoidance-Program/). The Traceability / Test Documentation Cover Sheet (Form 918) is always required, but the other documentation depends on the level of qualification met by the vendor. In general, the most qualified vendors (typically OCM/OEMs) require the least additional documentation while the least qualified vendors require more robust documentation for the traceability review. The traceability/test documentation requirements as listed on Form 918 are described below.

Traceability/Test Documentation Requirements in Accordance with DFARS 252.246-7008 & DLAD Procurement Note M01 Approved Suppliers for FSC 5961 Semiconductors and FSC 5962 Electronic Microcircuits:

(1) **Approved Source Manufacturer Specified in the Solicitation/Contract Item Description (Original Component Manufacturer (OCM), Original Equipment**

Manufacturer (OEM)) – Completed DLA L&M Form-918. If the bare item markings on the parts being provided do not match the CAGE and part number specified in the item description in the contract, the supplier must provide documentation certifying that the parts containing the alternate markings are the "exact product" as the item represented by the CAGE and part number specified in the contract item description.

(2) **Approved Source on the Applicable Qualified Products List (QPL)/Qualified Manufacturers List (QML)** – Completed DLA L&M Form-918. Also requires a Certificate of Conformance and Traceability (CoC/T), which must include information and documentation required by the applicable military specification for the QPL/QML product. Refer to DLA Technical and Quality Requirement RQ007.

(3) **Authorized Distributor of the OCM/OEM or QPL/QML Approved Source** – Completed DLA L&M Form-918 and evidence of a contractual arrangement with or the express written authority of the manufacturer or current design activity to buy, stock, sell, or distribute the part and an unbroken chain of traceability documentation through authorized distributors (if applicable), back to the approved source/manufacturer specified in the solicitation/contract.

(4) **Supplier/Distributor on the Qualified Supplier List of Distributors (QSLD) for FSC 5961/5962** – Completed DLA L&M Form-918 and evidence of an unbroken chain of traceability documentation, through trusted providers, back to an approved manufacturer specified in the solicitation/contract item description. Refer to the QSLD-5961/5962 document. QPL/QML items require a Certificate of Conformance and Traceability (CoC/T), which must include information and documentation required by the applicable military specification for the QPL/QML product. Refer to DLA Technical and Quality Requirement RQ007.

(5) **Supplier/Distributor on the Qualified Testing Suppliers List (QTSL) for FSC 5961/5962** – Completed DLA L&M Form-918 and complete test report including summary of test results, electrical testing read and record data, device photos, etc. Traceability documentation to the source of parts provided for testing. Refer to DLAD Procurement Note M01.

As previously described, the completed Form 918 and other documentation (as required) must be sent to the CDAP office at least fifteen days prior to the contract delivery date (CDD) via fax or email. This requirement allows the CDAP team sufficient time to review the documentation and provide approval or rejection notices as applicable. In the case where Form 918 has any errors, the CDAP team will notify the vendor and require resubmission of the documents. If Form 918 errors persist through more than two or three resubmissions the order is often cancelled.

If there are any discrepancies between the product number listed on the award and the product number provided by the vendor on Form 918, the CDAP team will first escalate the issue to a DLA Product Specialist or the appropriate Military Service Product Specialist before notifying the vendor. The Product Specialist will assess if the proposed substitute product is acceptable in lieu of the original requested part. If so, the alternate product number will be accepted, otherwise the substitute product will be rejected, and the order will be cancelled.

If all the information provided within Form-918 is acceptable, the CDAP team will review any additional documentation including test reports, traceability documentation, etc. This may include validating the chain of custody through traceability documentation, certification of the test report results, or other validation approaches. Once it has been determined that all submitted documentation is valid, the CDAP team will provide approval by signing Form 918 and sending it along with a ship letter to the vendor. This process also typically includes internal processes such as manually entering the Form 918 data into the CDAP catalog and uploading the documentation to Records Management (RM) within the Enterprise Business Systems (EBS).

1.3.3 Vendor Shipment

Upon receipt of a signed Form 918 from the CDAP team and approval to ship letter, the vendor is then able to complete the delivery process. The vendor must ship the materiel to the receiving warehouse co-located with CDAP at DLA Land & Maritime – Defense Depot Warren Ohio (DDWO) along with copies of all aforesaid required documentation – the signed Form 918, approval to ship letter, and any other required traceability documentation and/or test reports. These hard copies of documentation are required as further validation of the items throughout the CDAP process.

1.3.4 DDWO Receiving and Quick Look

When the shipped items and hard copies of required CDAP documentation are received at DDWO, the warehouse resources perform a process called a 'quick look'. This process involves comparing the information contained within the documentation to the parts themselves and the information entered in the CDAP catalog. If any information or product discrepancies are identified, a Supply Discrepancy Report (SDR) is generated. If the discrepancy is minor, the vendor is notified of a need to modify the order and resubmit all required documentation, otherwise the order is cancelled. However, if all provided information is valid and matches the items provided, the 'quick look' passes and is sent on to the test laboratory, which is also co-located at DLA Land & Maritime.

1.3.5 Product Test Laboratory Process

Once an FSC 5962 item is received at the test laboratory, the laboratory team performs a visual and x-ray inspection according to industry standards. If an item does not pass the visual inspection it is deemed unacceptable and the order is cancelled. If the part has significant defects or is suspected to be a potential counterfeit, the case is also sent for a legal review. However, if the item passes the visual inspection, it is put through a DNA marking process to help track the parts throughout the Department of Defense (DoD) supply chain. This DNA marking program is administered by the test lab and in conjunction with a DLA third party partner, Applied DNA Sciences (ADNAS).

1.3.6 Final Approval and Inventory

Once an item has passed a visual inspection and marked with DNA, it is deemed genuine the information for this contract is updated within RM to indicate the process has completed. The materiel is then passed back to DLA Procurement (J7) so that it could be added to the overall inventory and is eligible to be shipped to DLA customers.

1.4 Functional Pain Points to Mitigate

The existing CDAP processes include many opportunities for improvement and impact all stakeholders who participate in the process. Through requirements gathering and process mapping procedures, these problems were discussed with end users and documented to better understand how a future-state solution might address them and provide mutual benefit to all parties involved.

The functional pain points (or ‘problems’) identified during the analysis of these processes are depicted in Figure 2.

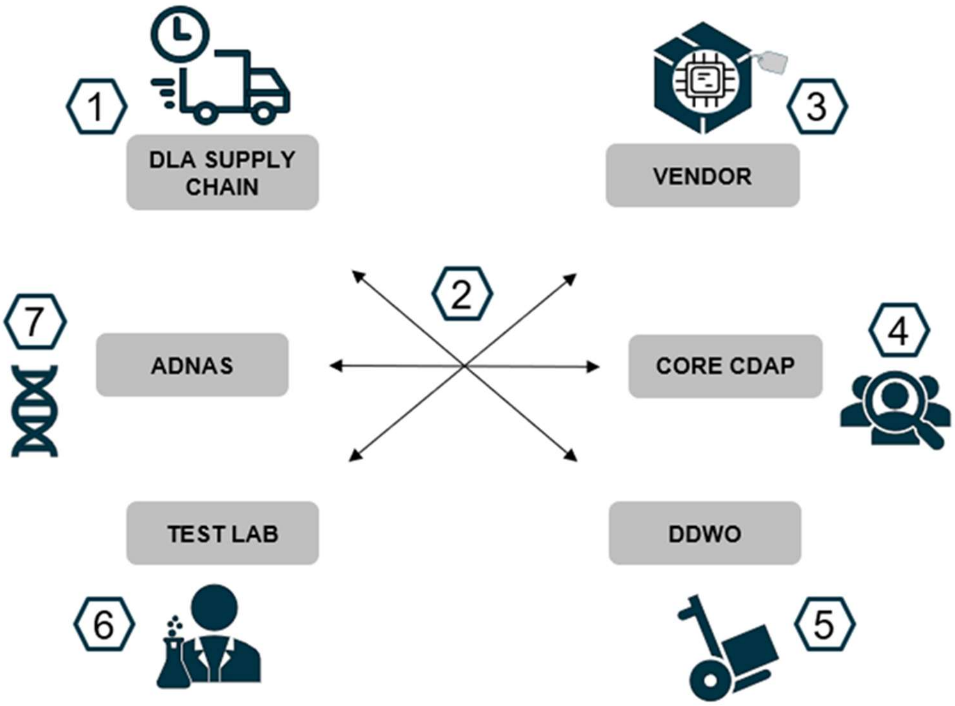

Figure 2. Functional Pain Points Throughout CDAP Processes

The functional pain points illustrated in the image above can be summarized as follows (in matching numerical order with the diagram):

1. **Error-Prone Process** – The existing process is slow and errors can trigger backorders. There is a need for a more automated process with reduced errors, so DLA can have improved material availability and reduced backorders for critical items.
2. **Lack of Visibility** – There is a lack of insight, efficiency, and traceability in the workflows between stakeholder groups. There is a need for more transparency and auditability.
3. **Redundant Manual Exchanges** – There is a redundant manual exchange of information between vendors and the CDAP Team. There is a need for enhanced credential sharing and data validation.
4. **Administrative Overhead** – There is a lack of time for the core CDAP team to focus on its primary role of reviewing traceability because of the administrative overhead and paperwork caused by errors.
5. **Lack of Shipping Insight** – There is a lack of insight into what will arrive at the DDWO warehouse and when. There is a need to establish better visibility of incoming procurements and their related shipments. There is a need to establish package traceability and sovereignty.
6. **Use of Many Databases** – There is a lack of workload awareness due in part to the use of multiple disassociated databases. Users within the Test Lab reference the CDAP Catalogue, their own database, and send data to ADNAS. There is a need for better awareness and data storage using a shared ledger.

7. **Data Entry** – The process includes redundant data extracts and time-consuming manual data entry by multiple stakeholder groups. There is a need for sharing data with the test lab and having auditability and tamper-evidence built-in.

1.5 Technical Pain Points to Mitigate

Just as there are multiple issues with the functional day-to-day CDAP processes, there are also improvements to be made to the existing technology employed by the program. Some of the major technical issues can be summarized as follows:

1. **Outdated Database** - The existing 'database of record' is built within Microsoft Access and will soon be phased out with no planned replacement.
2. **Lack of Enterprise Sharing** - There is no enterprise data sharing location for CDAP groups to track future work.
3. **No External Data Connection** - No options exist within DLA for sharing CDAP database information with a third-party partner.
4. **Lack of Automation** - Errors entering redundant data manually often lead to procurement delays and potentially backorders.
5. **Lack of Documentation Review Information** - It is difficult for vendors and CDAP groups to know the status of any given review, causing frustration and unnecessary email traffic.
6. **Lack of Data Re-Usability -** Nearly two-thirds of manually exchanged information is static and does not change on a frequent basis, however this information is often re-entered incorrectly by vendors and may lead to delays.

The technical issues listed above can be rectified in a number of ways. While the technical solutions for solving these issues may provide a variety of options, there are still some downsides to utilizing a traditional database approach instead of a blockchain ledger. Some of the considerations for potentially utilizing blockchain over these traditional approaches include:

- **Potential Data Tampering Issues** - Poorly implemented role-based access control (RBAC) could expose confidential data; moreover, the existing DB of record would still be centrally managed, and not readily tamper-evident.
- **Temporary Solutions** - Standard "lift and shift" migration to the cloud isn't cost-effective or future-proof.
- **Uncertainty Verifying External Data -** Digital credential exchange between manufacturers & distributors upstream could expose data and chain of custody data cannot be trusted as an auditable starting point for provenance.
- **Uncertainty Verifying Validity of Credentials** - With blockchain there is proof of non-revocation cryptographically (with non-repudiation). This allows for process automation for administrative lead-time reduction engendering a higher level of trust.

2.0 Blockchain Background

Blockchain is a new type of data system that maintains and records data in a way that allows multiple stakeholders to confidently share access to the same data and information. A blockchain is a type of distributed ledger technology (DLT), meaning it is a data ledger that is shared by multiple entities operating on a distributed network. This technology operates by recording and storing every transaction across the network in a cryptographically linked block structure that is replicated across network participants. Every time a new data block is created, it is appended to the end of the existing chain formed by all previous transactions, thus creating a chain of blocks called the blockchain. This blockchain format contains records of all transactions and data starting from the inception of that data structure. Blockchain is different from other types of data storage techniques in that transactions are immutable once added, meaning they are never edited or removed.

In most modern business applications where data is shared or passed between disparate entities or business partners, a one-way data flow exists where each entity maintains their own database and requires constant reconciliation with other stakeholders. In this type of architecture, each entity is the administrator of their database and makes decisions about whether (and when) to modify, append, or reject relevant data when it is received from an external partner. This very often results in multiple databases with differing information or differing states of data (i.e. not updated) even when they are intended to reconcile. Blockchains enable multiple stakeholders in an ecosystem to operate from a single shared set of mutualized data, eliminating the need for separate record keeping and reconciliation. In a blockchain construct, multiple parties can read and write to the distributed ledger while maintaining provenance, control, tamper evidence, and data integrity. The cryptographic approaches used in most blockchains provide a way to ensure data is verifiable; stakeholders can confirm transactions were added by other trusted stakeholders. These cryptographic approaches and key features of blockchain will be described later in depth. A depiction of key blockchain features is shown below.

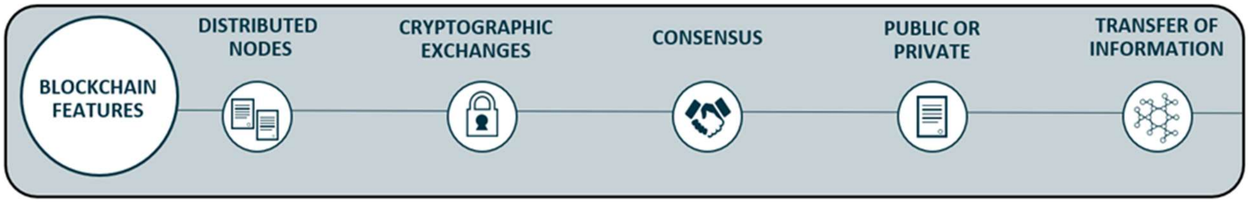

Figure 3. Depiction of Key Blockchain Attributes

For this prototype, we developed an **identity-based blockchain ledger** for storing vendor credentials and verifying them in real-time during the CDAP review processes. This means that the blockchain operating within our application is used to store new qualification credentials obtained by vendors from the list provided in the CDAP ***Pre-Award Qualification*** section. This approach ensures that CDAP (and Procurement users) could utilize the application to determine the validity of vendor credentials during the pre-award and post-award procurement processes, and enables a collaborative future-state approach with other government agencies where sharing of vendor qualification credentials is possible. Cybersecurity and cryptographic features of the blockchain will be described in subsequent sections and outline how data transfers across system nodes can be done in a secure and verifiable manner.

3.0 Cybersecurity

Cybersecurity is a key component of every Information Technology (IT) project, especially those that include the exchange of data packages over an interface or situations where DLA may be collaborating with external partners. In the case of the Counterfeit Detection and Avoidance Program, both conditions exist. Similarly, J6 and other Information Operations professionals may be familiar with the typical cybersecurity approaches required to certify new systems but may not be aware of how the use of a blockchain will impact those requirements or introduce vulnerabilities. Before commenting on cybersecurity considerations for this project and blockchain technology more broadly, it is important to highlight and understand DLA guidance on the topic as provided by the DLA Trusted Working Group with J6 representatives.

3.1 DLA Cybersecurity Guidance

The DLA-produced document ***Defense Logistics Agency Enterprise Enablers: Innovation, Data Management and Technology*** provides excerpts of guidance that reference cybersecurity directly, commenting on DLA Core Values and their alignment to cybersecurity reading as follows:

"We will continue our efforts in cybersecurity by being prepared for and able to adapt to changing conditions in the information environment. For example, we'll use virtual resources, managed with minimal resources and move more applications into the cloud with high availability. We will ensure services and capabilities are prepared for unplanned events with rapid failover to backup sites, simultaneously backing-up off-site, conducting regular continuity of operations exercises and increasing our cybersecurity focus on key mission areas. By investing in our future and ensuring support to the Warfighter, we further cement our importance and relevancy in logistics excellence." (quoted from page 3 of the document)

Additionally, a series of technology objectives are listed, among them is ***Technology Objective 3.1 – Maintain a Secure Environment***, which provides the following guidance:

"By using state-of-the-art technology, we will further our edge on information security and maintain a suite of cyber security hardware and software that is constantly monitoring and protecting DLA's information assets and people. On top of technology, DLA has already implemented to secure our information and the Warfighter we serve. These new advancements will position us to stand steps ahead of the enemy. The cyber threat is real and imminent and can be gravely damaging to DLA if not countered." (quoted from page 8 of the document)

Another document DLA published to outline clear guidance around the technical use of data & analytical tools is the ***DLA Data & Analytics Strategy*** document. This document provides a series of workstreams (called 'lines of effort') that guide considerations for various topics. Under ***Line of Effort 2: Security***, the following guidance is given:

"Security is the 2nd Line of Effort, which seeks to protect data from internal and external threats and from destructive actions through risk analysis, policies, standards, and technology. The CDO will assist in facilitating data security across DLA. Initiatives shall include the protection of data from accidental or intentional but unauthorized modification, destruction or disclosure using physical security, administrative controls, logical controls, and other safeguards to limit accessibility. Additionally, data security ensures that tools used across the IT environment meet cybersecurity requirements and standards. This is accomplished through controlled access to classified, proprietary, and privacy act data as well as implementation of role-based Public Key Infrastructure (PKI). Access to enterprise data assets will be controlled through the identification, authentication, and authorization of its business owners. Systems will use these access controls to determine access, viewing, modification or deletion of data assets. Data asset security will meet DOD and DLA security standards and be developed and reviewed by security stakeholders as part of the system development lifecycle." (quoted from page 8 of the document)

Finally, the ***DLA Strategic Plan*** published by DLA's Director provides overarching guidance for all enterprise operations. This plan outlines the Director's goals and objectives for the next several years, including guidance that references the need for improved Supply Chain Risk Management (SCRM), which may include the introduction of new technology such as blockchain. Under ***Objective 5.4: Mitigate Risks***, the document states,
[DLA should…]" Strengthen risk management to ensure secure, agile and resilient combat logistics support. DLA thoroughly manages risks associated with alternatives to deliver world-class logistics support. We must pay special attention to cyber risks and data integrity across the entire supply chain." (quoted from page 11 of the document)

3.2 Blockchain Cybersecurity Considerations

The guidance provided by DLA at numerous levels makes it clear that cybersecurity is a key focus area for all IT projects. However, how this guidance impacts the typical cybersecurity assessment practice may vary slightly for blockchain versus other more familiar technologies. As may be expected, a custom security architecture will be required for blockchain projects, depending on the type of distributed systems used and any identified gaps within the available platforms. As such, any guidance provided within this document should serve as a starting appoint for any and all cybersecurity assessments with the understanding that new and different cyber threats may be evolving. A full cybersecurity analysis should be conducted in partnership with relevant J6 departments during the pilot and production phases of any project with an increased focus on the to-be architecture as the project matures.

The fundamental cybersecurity principles of confidentiality, integrity, and availability always apply to blockchain implementations, as they would with other IT implementations. These core principles can be summarized as follows:

- **Confidentiality** – Are the systems and data kept secret or private from outside actors, and only made known to trusted and approved stakeholder groups?
- **Integrity** – Are the systems and data assured of their accuracy and devoid of influence or modification from outside actors?
- **Availability** – Are the systems and data available to be used when needed and not blocked, slowed, or obfuscated in any way by outside actors?

Blockchain based applications require security just like any other application, so the approach to meet these requirements does not change. However, the various risks, project context, and application characteristics may lead to a unique security design. In general, the security architecture required for a blockchain project will be different and more custom than usual. Some factors requiring a more custom security architecture include:

- **Blockchain is Distributed** – Distributed Ledger Technology is spread among multiple groups and potentially many organizational boundaries by definition. This may require alignment of these stakeholder groups and agreement on proper cybersecurity procedures and testing.

- **Blockchain Brings New Tools & Challenges** – Blockchain platforms and applications bring with them new untested tools and functional/technical challenges that must be addressed. By only allowing data to be added over time, and never removed or edited, blockchain data payloads may not work with certain systems or processes in ways that may not be immediately apparent. Fixed data variables, for example, will not 'update' on a blockchain since data can never be modified, so new approaches must be adopted, such as modifying the variable reference instead.
- **Blockchain Platforms are New** – Blockchain technology is expanding rapidly and many new software platforms are in a state of constant development to serve this demand. While new platforms serving specific use cases may provide a quick leg up in functionality and design, IT professionals must be deliberate in understanding the nuances around how data is handled, stored, and exchanged on different platforms. Some platforms naturally have capability gaps, and the underlying architectures required for these platforms may introduce cybersecurity gaps.
- **There is More Variance** – Blockchain architecture patterns seem to vary more than basic web applications or other well-known data structures such as relational databases. The variance in how software platforms architect their distributed nodes and systems means that each platform (and perhaps each implementation) will come with unique challenges to be understood and addressed.

Portions of the CDAP process that may include cyber touchpoints and an explanation of blockchain cybersecurity concepts are explained in later sections. However, some key aspects of utilizing blockchain and their related security considerations and processes are listed below.

Table 1. Blockchain Cybersecurity Components and Associated Security Activities

Cyber Component	**Blockchain**	**Security**
Governance	**Blockchain Ecosystem Rules and Permissions -** manage onboarding of participants into the network and the roles within the network.	• Identity and Access Management • Key Management via Hardware Security Modules • Information Security Guidelines and Policies • Security Training
Application	**Smart Contracts** are self-executing, logical workflow agreements between parties to manage access and use of value and data. **Third-party applications** can also be integrated within the platform.	• Application security • Secure code development and guidelines • Vulnerability assessment of third-party applications
Data	**On-chain Data Encryption** – Each data item can be encrypted individually, and	• Encryption and Key Management • Privacy policies and controls

Cyber Component	Blockchain	Security
	data can be segregated into on- and off-chain	• Industry specific standards and regulation • Off-chain Data Security
Transactions	**Blockchain Consensus and Rules** – Participants must agree on rules to commit transactions and to update the shared ledger with consensus algorithms.	• Secure and sustainable choice of consensus algorithm in design against double-spending, censorship • Fork management and maintenance
Infrastructure	**Blockchain Network & Systems** – Participants' nodes stand in their infrastructure and network, communicating through public or private connections	• Periodical Vulnerability Assessment and Penetration Testing • Logging and Monitoring • Endpoint Security and Patch Management

The activities listed above should help cybersecurity professionals understand which areas of a blockchain implementation should be included within pre-existing cybersecurity assessment activities, or where these activities should expand to account for the new technology. Understanding these associations will also help IT professionals identify gaps in the current assessment processes and where specific new analyses should address these concerns during a pilot phase or any time pre-production to set a cybersecurity baseline.

3.3 Blockchain Cryptographic Approaches

Blockchain is a relatively new technology, but it was created from the novel combination of pre-existing cryptographic approaches that date back decades in some instances. To best understand the cybersecurity implications of using blockchain technology, it may be useful to know what cryptographic approaches are commonly used with blockchain. The following sections will contain high-level summary explanations of complex technological components meant to convey the individual benefits contributed by each within a blockchain system. These explanations are taken directly from previous documentation produced leading up to this effort and provided in the ***Blockchain Overview*** document.

3.3.1 Asymmetric Key Encryption

Asymmetric key encryption, also known as ***public-key encryption,*** is a method for securely transmitting data using a pair of keys – a public key and a private key. The public key, which can be known by anyone and widely shared, is used to encrypt data being sent to the owner of the private key, resulting in an encrypted message as shown below.

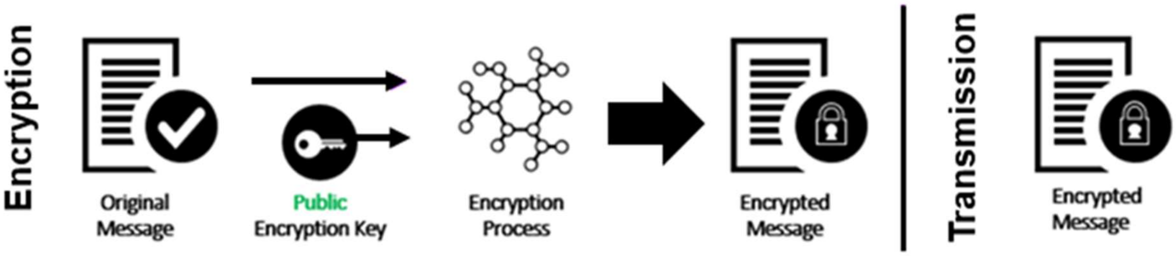

Figure 4. Encryption using Asymmetric Key Encryption

The private key, which is only known to the owner, can then be used to decrypt any message sent from holders of the public key, as shown in the figure below. Using this approach, encrypted data is unintelligible to anyone who intercepts the message during transmission without the closely-held private key.

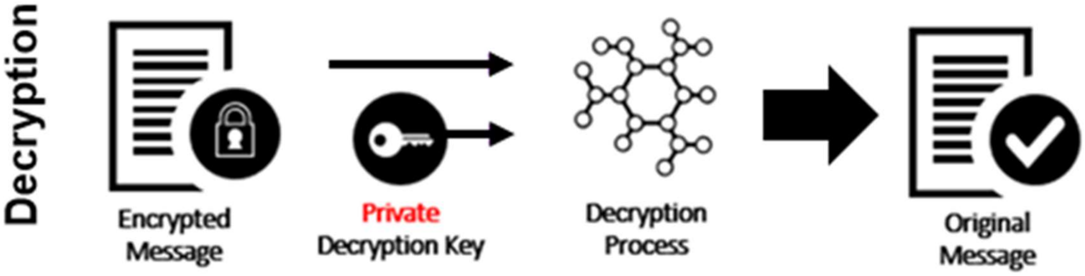

Figure 5. Decryption using Asymmetric Key Encryption

While encryption provides a secure method of transmitting data, it is not enough to ensure the long-term security of it. New encryption techniques are constantly in development, as well as malicious means to crack their processes over time. In response, new encryption standards are continually emerging to maintain effective guidelines. For these reasons, data segregation is the only reliable long-term strategy for ensuring the security of important and/or sensitive data.

Key Takeaway: Asymmetric key encryption enables secure transmission of data by utilizing public-private key pairs which are uniquely able to encrypt/decrypt data contents.

3.3.2 Hash Values

Hash values are alphanumeric values of a fixed length that uniquely represent an arbitrary amount of data and can verify the integrity of that data. When data (which can include any combination of text, images, etc.) is put through a hashing function, the output is a fixed-length string of alphanumeric characters known as the hash value. The most widely used cryptographically secure hash function is Secure Hashing Algorithm-256 (SHA-256), which creates a 256-bit hash value and was originally designed by the United States National Security Agency (NSA).

The power of cryptographically secure hash functions lies in two unique properties. First, hash functions are one-way functions, meaning that the original message contents cannot be reverse engineered from its corresponding hash value, nor can any information regarding the original message contents be discerned from a hash value. Second, a message passed through a given hashing function will always produce exactly the same hash value. If a single byte of the message content is altered, the resultant hash value changes. Thus, hash values provide a simple reference point for identifying whether a message being sent across a potentially unsecure channel has been altered. This is process is shown in the figure below.

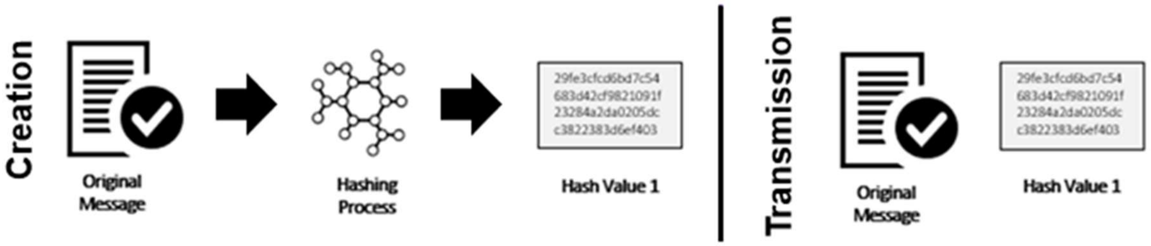

Figure 6. Hash Creation Process

A data package or message can be run through a specific hashing process (say SHA-256 for example) to create a hash value, which can then be sent to a recipient. The hash value alone is of no worth since it cannot be used to recreate the original message. However, the sender can then provide their original message to a recipient, who in turn can utilize the same hashing function as the sender to produce a hash value from the message as shown below.

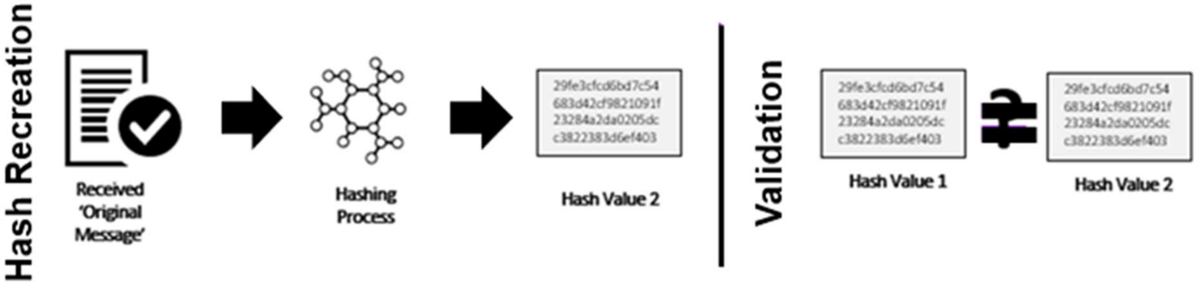

Figure 7. Hash Recreation and Validation Process

If the hash value that was originally sent matches this new hash value produced by running the message through a hashing function, then the recipient can be sure the contents of the message have not been altered in any way. This process is most useful when the sender does not care if the original message is intercepted and read, rather that they can confirm it has not been changed.

Key Takeaway: Hash values provide a method for ensuring data content has not been tampered with. This enables blockchains to be 'tamper evident'.

3.3.3 Digital Signatures

Digital signatures use a combination of asymmetric key encryption and hashing to both verify the identity of the signer and to verify that a document has not been altered, once signed. There are various digital signature algorithms, such as Rivest–Shamir–Adleman (RSA) or Elliptic Curve Digital Signature Algorithm (ECDSA). The inner workings of these algorithms are drastically different from each other, and the following graphic is an abstraction that is similar to RSA. As shown in the figure, this process is achieved by first hashing an original message and then 'digitally signing' that resultant hash by encrypting it using the sender's private key.

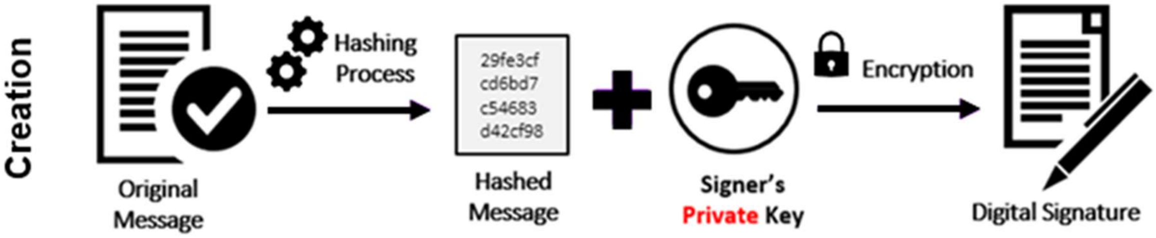

Figure 8. Digital Signature Creation Using Hashing and a Private Key

The signing process results in an original message which has been hashed and signed using encryption. The original message is then sent, unencrypted, to the recipient along with the digital signature. To confirm the authenticity of the file, a recipient can then take the message and run it through the same hashing process as the sender to create a hash value as shown in the figure below.

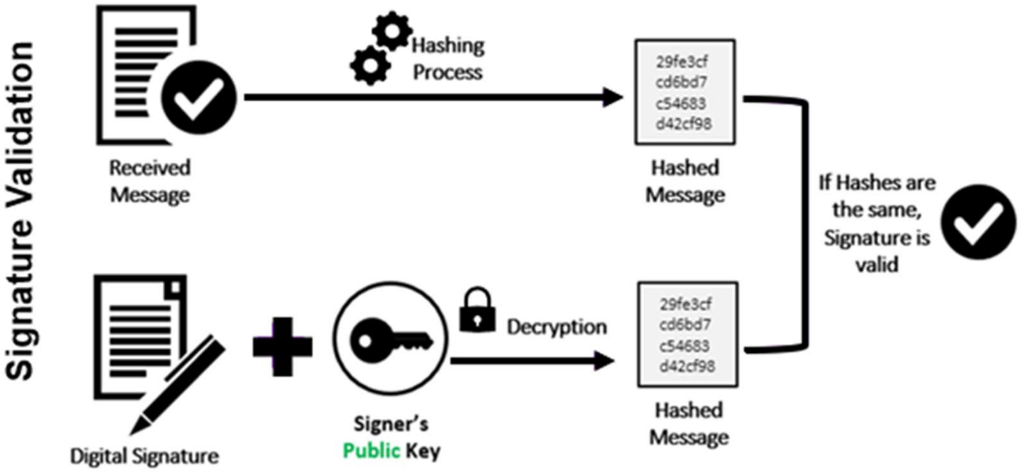

Figure 9. Process for Validating a Digital Signature Using Hashing and a Public Key

The recipient can then utilize the public key of the sender to decrypt the digital signature, resulting in another hash value. If the two hash values generated during these two signature validation steps are identical, then the authenticity of the sender is confirmed

Key Takeaway: Digital Signatures utilize hashing and asymmetric encryption to enable authentication of a sender's identity.

3.3.4 Merkle Trees

Merkle Trees are trees of hashes that allow sets of data to be verified using only portions of the tree. Their basic structure consists of a 'starting' root block upon which subsequent blocks are appended in sequential order, retaining hash value information from the previous block. In traditional databases, data is entered (grouped) and constantly altered. This method creates two problems. First, if a database contains too much information, it can become computationally arduous to interact with or to protect (via hashing) the data. Second, the data in traditional ledgers or databases can usually be altered by anyone with access. This can be problematic, because unwanted or malicious changes are difficult to track and trace, specially finding who was responsible for said changes. Merkle Trees present a solution to both problems. Each piece of data on a single block receives a hash value that is connected to a root hash as shown in the figure below.

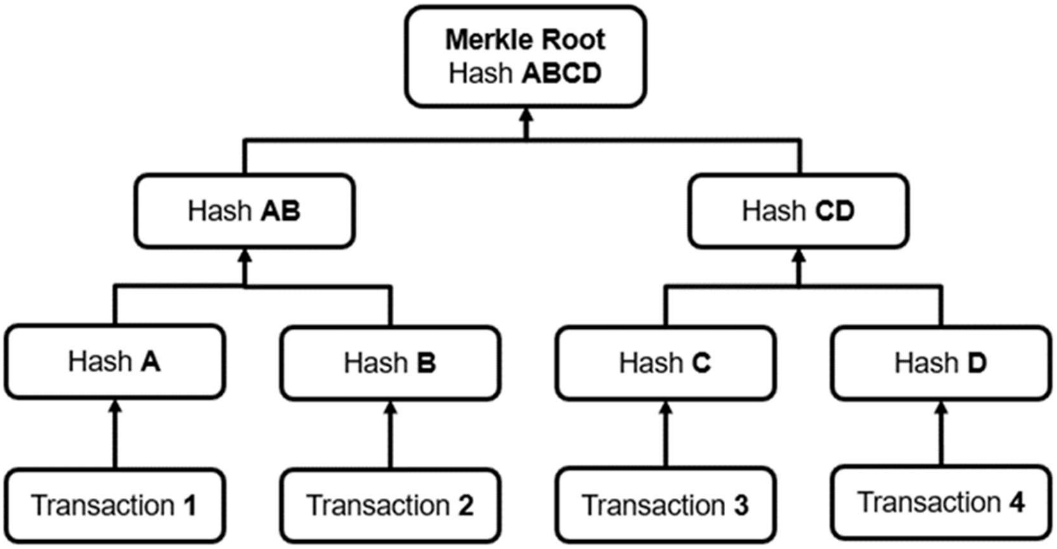

Figure 10. Merkle Tree Representation of Four Transactions on a Single Block

This method requires less computational power, because each transaction or data entered is hashed separately as compared to traditional systems where data is grouped and there is continuously hashing as data is entered or edited. The other solution Merkle Trees offer is data immutability and auditability. This means that once data is entered it cannot be changed, and everything is traceable as once a transaction is made a signature is attached to it. If there is any alteration in the blockchain, it will alert all the users and it will not compute as the hash values in the 'tree' will not match with the original values. Merkle Trees can be considered an optimization component, as they help avoid huge transaction/block sizes and computational overheads. It is important to note that Merkle Trees are not used for all platforms (e.g. Corda).

Key Takeaway: Merkle Trees are a tree of hashes that allow sets of data to be verified using only portions of the tree.

3.3.5 Peer-to-Peer Networks

Peer-to-Peer networks are those in which participants are communicating and exchanging data directly with each other, without the need for a central authority to control data flows. It is this type of network setup which allows nodes within a blockchain ecosystem to share data and validate datasets with each other directly. When consensus is reached, and a new block is added to a blockchain, the participants within a peer-to-peer network communicate with each other and ensure that their new copies match the new copies of every other participating node. It is by communicating directly with each other to compare data that participant nodes can achieve consensus. Peer-to-peer networks provide a means for public blockchains to add any number of participants and ensure that blockchain copies are widely replicated and geographically distributed.

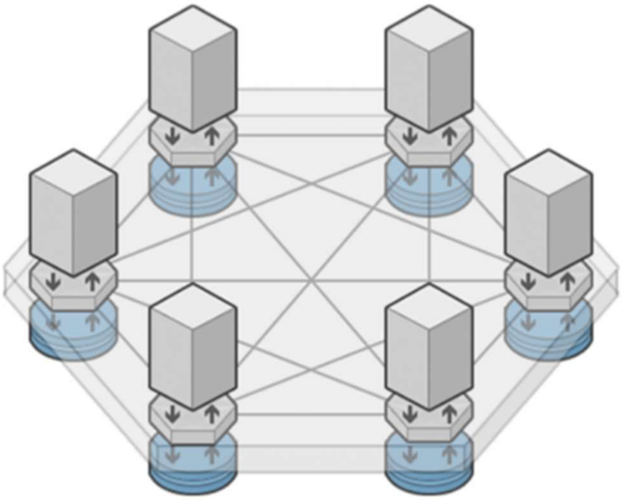

Figure 11. Diagram of Peer-to-Peer Network

A large number of cooperating nodes within a network makes it increasingly difficult for a malicious actor to sabotage the network, since they would need to alter or disrupt an increasingly large number of nodes to do so. The additional redundancy also helps ensure that the blockchain will continue to operate in case some subset of nodes fail, malfunction, or become inactive for any reason.

Key Takeaway: Peer-to-Peer Networks are those in which participants are communicating and exchanging data directly with each other, without the need for a central authority.

3.3.6 Smart Contracts

A ***smart contract*** is a component of a blockchain-based system that can automatically enforce participant-agreed rules and process steps. In traditional database terms, smart contracts are similar to ‘stored procedures’, but in this case represent an agreement between multiple entities. Smart contracts are coded business rules which can facilitate, verify, and execute the terms of an agreement between counterparties without the need for a human intermediary. Once launched, smart contracts can be either fully autonomous or triggered actions. When contract conditions are met, pre-specified and agreed actions occur automatically. In the supply chain realm, smart contracts could enable inventory orders, predictive and automated maintenance, payments, and information transfer. The use cases begin to expand greatly when additional emerging technologies are integrated, for example when Internet of Things (IoT) devices are utilized as sensory input for smart contract conditions. Insurance agreements are one other example of how smart contracts can be used in the future to simplify the claims process.

For example, insurers can set specific conditions for paying out insurance claims based on any type of available data (e.g. if no rain falls on a property for two consecutive weeks and the temperature exceeds ninety-five degrees, then a drought insurance policy will pay out). Smart contracts enable the logic-based exchange of any digitally-stored assets, making it possible to perform common transactions such as making insurance payments, vehicle purchases, real estate purchases, or stock exchanges without the need for an intermediary, and at a reduced cost.

Key Takeaway: Smart contracts are an added feature of some blockchains which allow scripting languages and the ability to execute coded logic to be included in blockchain structures.

3.4 Technical Architecture Decisions Impacting Cybersecurity

Cybersecurity should be considering during all phases of design and implementation of a blockchain solution. The cybersecurity considerations described previously were used to set forth a set of principles and design decisions impacting the overall technical architecture for this project.

The major cybersecurity architecture design decisions in place for this project are as follows:

- No data stored on-chain should contain private, confidential, or personally identifiable information (PII).
- The chosen cloud platform must be certified by the Federal Risk and Authorization Management Program (FedRAMP).
- Application components must be encapsulated into containerized microservices.
- Open data standards of Verifiable Credentials 1.0 and Decentralized Identifiers 1.0 developed by the World Wide Web Consortium (W3C) must be used.
- The prototype must ensure data stored in the system is tamper-evident.
- Blockchain capabilities must be available through application programming interfaces (APIs).
- The prototype must have the capability to grant, validate, and revoke access across elements of the system via Role-Based Access Control (RBAC).
- The prototype should be able to authenticate users and participants of the network.
- The prototype should define and support a mechanism to manage identities.
- The prototype must be able to create cryptographic proofs for the verification of credential claims such as its attributes, non-revocation, the credential's holder, and the credential's issuer.
- The prototype should allow for the revocation of issued credentials.

3.5 Cybersecurity Touchpoints in CDAP

The future-state CDAP processes will include digital exchanges of information with internal stakeholders and external vendors. Additionally, the background cryptographic concepts explained in previous sections will be utilized to provide different data assurances throughout the process. The major points within the process where cybersecurity and cryptography are concerned include the following (labeled on the figure on the next page):

1. **Vendor Sends Form 918 Information, Qualification Credential, and Additional Documentation to CDAP**
    - **a.** DLA can attest that the vendor is qualified according to their 'Qualification Credential' level by verifying the digital version of the credential stored on the blockchain.
    - **b.** The verification process includes (i) verifying the credential holder's digital signature against the vendor's, (ii) verifying DLA is the entity that issued the credential, (iii) verifying the credential has not been revoked, and (iv) verifying that the qualification stated on the credential is aligned with the awarded contract CDAP requirements.
    - **c.** Vendors providing parts will exchange demographic information, traceability documentation and test lab results via a communication channel encrypted end-to-end using authentication encryption methods.

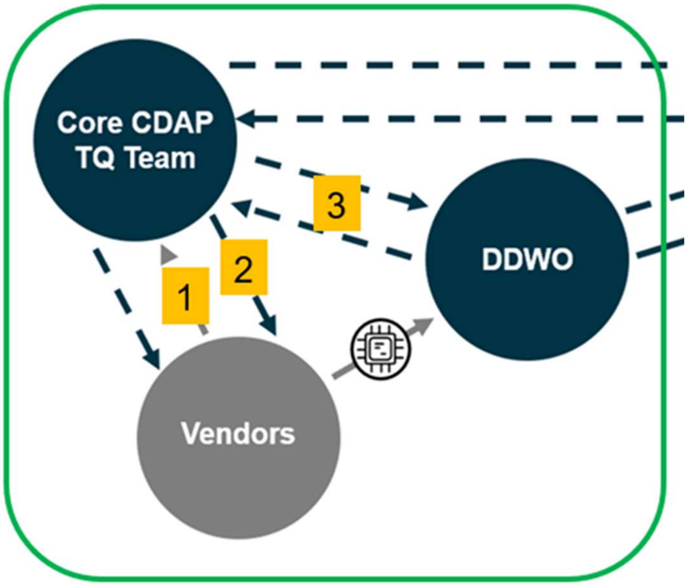

Figure 12. Cybersecurity Touchpoints within CDAP Process

2. **CDAP Reviews Form 918 Information, Qualification Credential, and Additional Documentation before sending Authorization to Ship**
   a. The Technical Quality (TQ) team can execute CDAP process steps digitally via a web application interface that is able to establish a peer-to-peer connection with the vendor, allowing the vendor to exchange supporting documentation with DLA TQ faster and securely.
   b. A barcode or Quick Response (QR) code can be used instead of adding a hard copy of the Form 918 in the shipment packaging to facilitate its verification once the shipment arrives to the DLA Warehouse (DDWO). This type of barcode creation and scanning would be completed in a pilot phase.
3. **DDWO Confirms Form 918 Information matches the Information on file and Performs First Look**
   a. DDWO can verify the shipment by cross-checking it against what it is stored on the blockchain. Verification is achieved by scanning the barcode or QR code on the box. This barcode or QR code contains the hash value of the information submitted in the Form 918. If the hash exists on the blockchain ledger, then the form is valid. If not, the shipment will be automatically flagged for further investigation. This type of physical scanning process can be demonstrated during a pilot phase.

4.0 Governance Approach

As a type of distributed technology, one key aspect of all blockchain implementations is the governance approach – an agreed upon set of rules, procedures, and processes for making decisions that impact the overall blockchain ecosystem of partners. In cases such as this prototype where DLA is the primary governing body and no other partner systems are impacted, DLA can make unilateral decisions. However, it is in DLA's best interest to plan and account for the final distributed end-state use of the blockchain and potential technical impacts when deciding how to approach governance. Additional governance guidance from the ***Blockchain Overview*** document is presented below with project-specific guidance woven in where applicable.

4.1 Type of Permissions

One of the initial governance decisions needed is defining the type of blockchain that will be used. Depending on the degree of privacy required, different types of blockchain exist. Finding the right type is key and depends on the purpose of the respective use case. In the most basic sense, there are two types of blockchains to consider: ***permissioned*** and ***permissionless***.

In a ***permissionless*** blockchain, anyone can have access to the underlying data and transaction history. Furthermore, all participants in the network are treated as equal, meaning that all users have equal rights to read data and execute transactions. Most cryptocurrencies (e.g. Bitcoin) operate on permissionless blockchains. They are frictionless for anyone to transact on and provide everyone the ability to access a complete copy of the transaction history (ledger).

In a ***permissioned*** blockchain, one or more organizations control who can have access to the underlying data and transaction history. Additionally, user identities are authenticated and known through some type of procedure (e.g. Know Your Customer). Different levels of read and write access can be assigned to participants for various types of data in the distributed ledger. Permissioned blockchains are often appropriate for enterprise use cases within sectors such as supply chain, healthcare, financial services, etc. The onboarding of ecosystem participants is generally controlled by one or more organizations and user identity is authenticated and known. Because of this, we can assign different levels of read and write access to ecosystem participants, for various types of data in the blockchain. This provides greater control and privacy considerations to an enterprise than permissionless blockchains.

As a system fully governed and operated by DLA, this prototype was built as a ***permissioned*** blockchain. However, even if this system was operated in conjunction with government and external partners owning nodes, this system would remain permissioned. In nearly every conceivable private business (or government) use case, the underlying structure should remain permissioned, where access to the system is closely controlled and set in agreed upon governance rules.

4.2 Key Governance Roles

The foundation for successful blockchain initiatives is less about a technical implementation and more about building a willing and cooperative ecosystem of participants with well-established roles. Governance establishes a process to achieve agreement across each organization in the ecosystem through a formal set of processes, communication forums, roles, responsibilities, and tools. It is important to establish and automate the governance processes so that they are known and agreed to by all ecosystem partners, with the opportunity for them to be encoded into smart contracts on a blockchain network if that may provide a functional or technical benefit. Common important roles required for each ecosystem of blockchain partners are described below.

Table 2. Framework for Governance Prioritization within a Partner Ecosystem

Governance Role	Description	Priority Activities	Value
Operationalize Ecosystem	Create and operate a common set of rules including: • Solution • Products & Services • Standards & Rules • Governance • Partners & Membership • Operations • Business Model	• Drive founding member engagement and coordination • Develop and manage plan for ecosystem to move forward with testing and enabling transactions • Define ecosystem roadmap	High
Project Management Office	End-to-end management of the design, build, and launch of the ecosystem including the development of the common set of rules.	• Important for oversight and execution • Coordination amongst various stakeholders	High
Provide Platform & Other Services	Build and operate blockchain services (if required) including reporting, exception handling, validation, routing, dispute management, etc.	• Critical services need to be designed, enabled, and optimized to scale the solution & ecosystem	High
Ecosystem Advisor	Provide best practices, insight and guidance to establish, manage and optimize the ecosystem.	• Improve speed to operation and collaboration amongst ecosystem members.	Medium

4.3 Other Governance Considerations

While designing a governance model, it is important to consider the set of standards and rules that will apply to day-to-day governance of the blockchain network, as well as a business model and risk management services. In most blockchain governance models, there are several layers of governance, such as: entity, organizational, leadership and data governance. Below are several the key considerations an ecosystem should consider when establishing their governance model, with commentary on how these issues were approached for this project and future considerations:

- **Network Governance Structure –** In the prototype phase, DLA owns the governance and decision-making capabilities for the entire network. If this, or any other blockchain capability moves into a pilot or production capability with other government partners a shared governance model should be adopted.
- **Defined Roles and Responsibilities for all Ecosystem Partners –** Roles and responsibilities can be divided any way the stakeholder group decides. However, it is important that all roles are filled and that any important functional or technical responsibilities are understood. For example, in this prototype phase DLA owned and operated all nodes of the blockchain, hosting virtual instances for the external vendors. This approach ensures a higher level of security, but also removes some of the benefits of having a decentralized system. In cases where other government partners may want to participate in the ecosystem, it would make sense to allow them to own and operate their own network nodes.
- **Common Data Standards –** The primary use of a blockchain structure is to store and exchange data. As such, it is extremely important that data payloads are structured in an agreed upon manner. This is similar to many existing systems where strict data controls are required and was prototyped in this phase by providing field-level validations and controls on user input forms. Again, this consideration may become more important when major government partners wish to share similar data with DLA.
- **Interoperability with Existing Data and Standards –** Interoperability describes the ability of a blockchain or other data connection (such as an interface) to seamlessly send and/or receive data into (or out of) another system or database. Many systems may operate incorrectly or unexpectedly if data is received or sent in an unplanned format, so this is important to consider. Blockchain ledgers store data payloads in ways that may be unfamiliar or unique to their platform, so it is important to understand the way data is sent to the blockchain or retrieved from it in all data exchange steps.
- **Process to Update Rules –** Once rules for operation have been established among an ecosystem, they may be difficult to change, especially if those rules relate to modifying the storage of data or system inputs/outputs in general. As such, it is important to consider procedures for updating, adding, or removing rules if such a need should arise. All ecosystem partners should agree to this process before the blockchain system goes live so that established rules will remain objective to all parties.
- **Onboarding and Offboarding –** All partners must agree upon a standardized process for onboarding new stakeholder groups (if applicable) and individual user accounts. During this prototype, it was determined that all stakeholder groups (including external vendors) would be responsible for onboarding their own individual user accounts. However, for

obvious reasons the Core CDAP Team is the only group that is allowed to grant access to new stakeholder groups and vendor company-level accounts.

- **Decentralized Key Management Solution(s)** – Encryption keys are paramount to cybersecurity and the successful operation of a blockchain network. As such, it is very important to agree upon the most secure method for managing these keys and where/how they should be stored. This will largely depend upon the type of solution being developed and the software used to build the capability.
- **Security Control** – Security is a very broad term, but cybersecurity roles should be determined and agreed upon by all stakeholder groups. In some cases, this may simply include agreeing upon a set of shared minimum cybersecurity standards but may also include activities such as periodic penetration testing, resiliency, and cybersecurity certification.

4.4 Prototype & Pilot Ecosystems

Within this prototype build, the systems used to operate the application were completely 'owned' by DLA and virtual nodes were stood up for vendor partners. This primarily means that the instances were managed and controlled by DLA, without an apparent downside for the vendors. This allowed the enterprise to focus on key aspects of prototyping the capability such as designing a mutually beneficial functional solution, determining the value of the solution, gaining buy-in from stakeholders, and understanding the implications of high-level governance decisions.

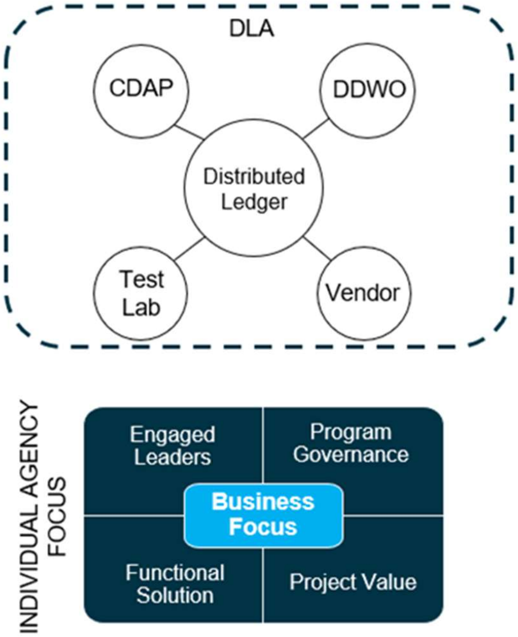

Figure 13. Prototype Ecosystem - Distributed Nodes Contained Within DLA

For the potential blockchain pilot phase, a different ecosystem approach would need to be explored, including simulation of distributed nodes owned by each of the stakeholder partners. This expansion of the ecosystem may also include government partners and would benefit from demonstrating and testing the technical implications of distributing the nodes. As described above, this approach may also make certain unique governance considerations apparent that may otherwise go unnoticed. The proposed distributed ecosystem with a minimum of three partner nodes (minimum viable ecosystem or MVE) is shown on the following figure.

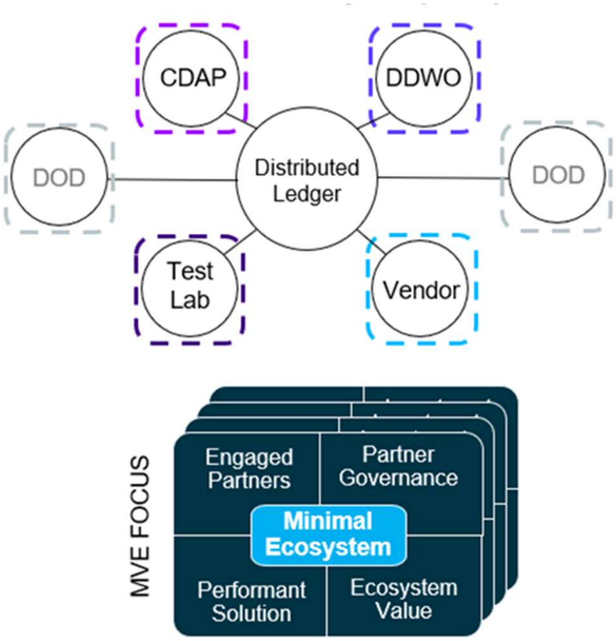

Figure 14. Potential Pilot Ecosystem – Distributed Nodes within Participant Domains

The expanded approach recommended for the pilot phase would ensure that multiple trusted partner organizations would have the opportunity to ensure functional and technical success. This includes the proper functioning of the system during normal use (functional operation), demonstration of the system value, and any technical and/or security concerns that may become apparent.

5.0 Requirements & Functional Design

One key success factor for blockchain implementations is the development of a collaborative environment that is mutually beneficial for all parties involved. As such, it was critical that functional requirements describing how the system will function and the features it will include were collected from all stakeholder groups including both internal and external entities.

5.1 High-Level Requirements

The requirements gathering process was completed by interviewing each stakeholder group separately and documenting their needs for a new capability. All of these ideas and features were collected together and combined to form the 'Recommended Requirements' for the project, which would form the basis of the functional design. Ultimately, the detailed requirements for this project fell under eight major high-level requirements as described below.

1. **User Roles -** The types of User Roles available shall be dictated by User Group.
2. **User Role Capabilities -** User Roles will dictate the functional controls and capabilities of each User Account.
3. **CDAP Registration Process -** A CDAP Registration Process will be put in place to facilitate the creation of new Organizations within the Counterfeit Detection and Avoidance Program ecosystem.
4. **User Account Creation -** A User Account Registration Process will be put in place to facilitate the creation of new User Accounts within a specific Company/Organization.
5. **Credentialing Process -** A Credentialing Process will be put in place to facilitate the attestation of Qualification Statuses to External Organizations.
6. **Post-Award Documentation -** A process shall be developed for vendors to submit documentation tied to an award for review and approval by the CDAP team.
7. **Form 918 Information -** A process shall be developed for vendors to submit required Form 918 Information via a web-based Form.
8. **General Web Form Features -** All web-based forms will include certain best practice features.

5.2 Detailed Requirements

Each of the major high-level requirements contains several detailed-level requirements describing the desired functionality, fields, and other factors required to meet the Trusted Working Group's intended end-state functionality. The detailed requirements are listed below corresponding to their high-level requirement groupings.

Table 3. Requirements under High-Level Requirement 1: User Roles

HLR-1	The types of User Roles available shall be dictated by User Group
FR-1-1	User Role types will include the following four roles at a minimum: Universal Administrator, Administrator, Standard User, and Auditor
FR-1-2	The Core CDAP team shall be the only User Group with access to universal administrator roles
FR-1-3	The Core CDAP Team, Internal Organizations, and External Organizations shall have access to Administrator, Standard User, and Auditor User Roles

Table 4. Requirements under High-Level Requirement 2: User Role Capabilities

HLR-2	User Roles will dictate the functional controls and capabilities of each User Account
FR-2-1	User Roles will dictate the functional controls and capabilities of each User Account as defined in the User Role Capability Matrix
FR-2-2	All User Roles except Universal Administrator shall define the functional controls and capabilities of each User Account within a specified Organization only
FR-2-3	User Roles will define database access and read/write capabilities for each User Account
FR-2-4	User Roles will define document schema creation/modification capabilities for each User Account
FR-2-5	User Roles will define document submission and approval capabilities for each User Account
FR-2-6	User Roles will define credential schema creation/modification capabilities for each User Account
FR-2-7	User Roles will define credential attestation capabilities for each User Account
FR-2-8	User Roles will define onboarding capabilities for each User Account
FR-2-9	Only Administrators can Modify Company/Organization Demographic Information

Table 5. Requirements under High-Level Requirement 3: CDAP Registration Process

HLR-3	**A CDAP Registration Process will be put in place to facilitate the creation of new Organizations within the Counterfeit Detection and Avoidance Program ecosystem**
FR-3-1	The CDAP Registration Process will be facilitated via a web-based form
FR-3-2	The CDAP Registration form will have the following data fields at a minimum: Applicant First Name, Applicant Last Name, Applicant Email Address, Applicant Phone Number, Company/Organization Name, Company/Organization Physical Address, Company/Organization CAGE, Desired Qualification Type, Additional Documentation (Attachment Option), Additional Information (Text Box)
FR-3-3	All fields within the CDAP Registration form will be required for submission except for Additional Documentation (Attachment Option) and Additional Information (Text Box)
FR-3-4	Administrators within the Core CDAP Team will be able to review submitted CDAP Registration Forms and approve or reject the applications
FR-3-5	Upon Approval of a CDAP Registration Form, the Company/Organization will be created within the CDAP ecosystem
FR-3-6	Upon Approval of a CDAP Registration form, the Applicant will be issued a user account using the email address and password submitted on the application form.
FR-3-7	Upon Approval of a CDAP Registration Form, the Applicant User Account will be given the User Role of Administrator
FR-3-8	Applicants will be notified at the Applicant Email Address when a CDAP Registration Form application has been approved or rejected.

Table 6. Requirements under High-Level Requirement 4: User Account Registration

HLR-4	**A User Account Registration Process will be put in place to facilitate the creation of new User Accounts within a specific Company/Organization**
FR-4-1	The User Account Registration process will be facilitated via a web-based Form
FR-4-2	The User Account Registration Form will have the following data fields at a minimum: Applicant First Name, Applicant Last Name, Applicant Email Address, Applicant Phone Number, Company/Organization, Company/Organization CAGE, Additional Documentation (Attachment Option), Additional Information (Text Box)
FR-4-3	All fields within the User Account Registration Form will be required for submission except for Additional Documentation (Attachment Option) and Additional Information (Text Box)
FR-4-4	Administrators within the Company/Organization listed on the User Account Registration Form will be able to review submitted User Account Registration Forms and approve or reject the applications
FR-4-5	Submitted User Account Registration Forms will be routed to the appropriate Company/Organization based on an exact match of the Company/Organization CAGE as submitted by the applicant
FR-4-6	Upon Approval of a User Account Registration Form, the approving Company/Organization Administrator will assign the new account a role type based on available roles.
FR-4-7	Upon Approval of a User Account Registration Form, the Applicant will be issued a user account using the email address and password submitted on the application form.
FR-4-8	Applicants will be notified at the Applicant Email Address when a User Registration Form application has been approved or rejected.
FR-4-9	A process shall be put in place for resetting forgotten passwords

Table 7. Requirements under High-Level Requirement 5: Credentialing Process

HLR-5	**A Credentialing Process will be put in place to facilitate the attestation of Qualification Statuses to External Organizations**
FR-5-1	Each Organization/Company registered through the CDAP Registration Process shall have the ability to create their own Credential Schemas and assign Credentials to External Organizations (i.e. any Organization outside of theirs) as deemed appropriate
FR-5-2	Whenever an Organization/Company creates a new Credential schema, this Credential is available for assignment to other External Organizations and/or User Accounts within the assigning company's ledger
FR-5-3	An Organization/Company can only be assigned a Credential if documentation is attached to the record as proof of their qualification.
FR-5-4	Applications to the CDAP Credentialing Process can be submitted via web-based Form
FR-5-5	The CDAP Credentialing Process application Form will have the following fields at a minimum: Applicant First Name, Applicant Last Name, Organization/Company Name, Organization/Company CAGE, Desired Qualification Type, Qualification Documentation (Attachment Option) , Additional Information (Text Box)

FR-5-6	All fields within the CDAP Credentialing Process application Form will be required for submission except for Additional Information (Text Box)
FR-5-7	The Core CDAP Team shall have the ability to assign the following 5 credentials (at a minimum): (1) Approved Original Component Manufacturer/Original Equipment Manufacturer, (2) Approved Source on the Applicable Qualified Products List/Qualified Manufacturers List, (3) Authorized Distributor of the OCM/OEM or QPL/QML Approved Source, (4) Supplier/Distributor on the Qualified Supplier List of Distributors for FSC 5961/5962, (5) Supplier/Distributor on the Qualified Testing Supplier List for FSC 5961/5962
FR-5-8	Any Organization issuing a credential has the ability to revoke the credential at any time

Table 8. Requirements under High-Level Requirement 6: Post-Award Documentation

HLR-6	**A process shall be developed for vendors to submit documentation tied to an award for review and approval by the CDAP team.**
FR-6-1	Credentialed Vendors will receive a notification when a new FSC 5962 award has been made to their registered CAGE
FR-6-2	Credentialed Vendors can submit documentation for review tied to a specific award/purchase order
FR-6-3	Credentialed Vendors shall be able to review the status of any previously submitted documentation
FR-6-4	The Core CDAP team can review documentation submitted for an award and provide comments back to the awardee (Credentialed Vendor)
FR-6-5	The Core CDAP team can review documentation submitted for an award and approve or reject the documentation
FR-6-6	The Core CDAP team can provide a Credentialed Vendor with approval to ship for a specific award
FR-6-7	Credentialed vendors will receive a notification when documentation has been approved or rejected
FR-6-8	Credentialed Vendors shall have the ability to indicate whether a document submission is for a ‘partial shipment’ and allow the record to ‘stay open’ for additional document submissions and reviews if this is the case.
FR-6-9	Upon documentation submission, the system shall check the status of the vendor’s attached credential and only send through submissions if the credential is active and verified.

Table 9. Requirements under High-Level Requirement 7: Form 918 Information

HLR-7	**A process shall be developed for vendors to submit required Form 918 Information via a web-based Form**
FR-7-1	Demographic Information within Form 918 will be pre-populated based on the Company/Organization tied to the User Account submitting the form
FR-7-2	The web-based Form 918 will include a Supplier Information section with the following fields auto-populated based on stored information about the external vendor: a. Supplier Name b. Supplier Address i. Street ii. City iii. State iv. Country v. ZIP code c. Commercial and Government Entity (CAGE) code d. Point of Contact / Submitter Name e. Point of Contact / Submitter Email f. Point of Contact / Submitter Phone Number
FR-7-3	The web-based Form 918 will include a Submitter Information section with the following fields auto-populated based on stored information about the person submitting the form: a. First Name b. Last Name c. Email d. Phone Number
FR-7-4	The web-based Form 918 will include a Contract Information section with the following fields auto-populated based on stored information about the contract: a. Contract Number b. Contract Line Item Number (CLIN)
FR-7-5	The web-based Form 918 will include an Item Information section with the following fields auto-populated based on stored information about the item: a. Item Name b. Federal Supply Classification (FSC) code c. National Item Identification Number (NIIN)
FR-7-6	The web-based Form 918 will provide a method for indicating whether an external vendor has changed their Company Name or Company Address.
FR-7-7	The web-based Form 918 will include a Quantity / Date Code section requiring the following information about the items: a. Date Code b. Quantity
FR-7-8	The web-based Form 918 will provide a method for entering multiple unique Date Code and Quantity combinations.

FR-7-9	The web-based Form 918 will provide a Traceability Documentation Type including a method for the external vendor to certify their level of qualification with the following options: a. Approved Source Manufacturer Specified in the Solicitation/Contract Item Description (Original Component Manufacturer (OCM), Original Equipment Manufacturer (OEM), b. Approved Source on the Applicable Qualified Products List (QPL) / Qualified Manufacturers List (QML), c. Authorized Distributor of the OCM/OEM or QPL/QML Approved Source, d. Supplier Distributor on the Qualified Suppliers List of Distributors (QSLD), e. Supplier/Distributor on the Qualified Testing Suppliers List (QTSL) for FSC 5961/5962
FR-7-10	The web-based Form 918 will provide a method for viewing information / help text about each of the qualification levels (i.e. QSLD, QSLM, QML, etc.) and what types of Traceability Documentation are required for each
FR-7-11	The web-based Form 918 will include a Type of Documentation section including a method for the user to select the type of Traceability Documentation they are submitting from the following choices: a. Traceability Document(s) b. Test Report(s)
FR-7-12	The web-based Form 918 will provide a method for submitting Traceability Documentation files
FR-7-13	The web-based Form 918 will require a user to digitally sign the form and enter the date before confirming their information and submitting it
FR-7-14	The web-based Form 918 will include an Item Information section requiring the following information about the item: a. Part Number b. Original Part Manufacturer c. Manufacturer CAGE Code
FR 7-15	If the external vendor indicates they have changed their Company Name or Company Address, the CDAP team will be sent notification and must provide approval before the changes take place within the system.

Table 10. Requirements under High-Level Requirement 8: General Web Form Features

HLR-8	**All web-based forms will include certain best practice features**
FR-8-1	The web-based Forms will be compatible with Google Chrome and/or other standard web browsers as determined by J6 and the stakeholder group.
FR-8-2	The web-based Forms shall provide information buttons and/or another means for conveying information about completion of the fields to users
FR-8-3	The web-based Forms will be accessible to external users without the use of a common access card (CAC)
FR-8-4	The web-based Forms will include built-in validation rules to ensure proper data formats for certain fields prior to submission and according to field characteristics in the data dictionary, as applicable
FR-8-5	The web-based Forms will provide a method for submitting multiple Documentation files
FR-8-6	The web-based Forms will provide a method for changing or deleting previously submitted Documentation
FR-8-7	The web-based Forms will provide a summarized overview of all submitted information for user confirmation before being final submission
FR-8-8	The web-based Forms will provide a method for submitting the information (i.e. a submit button)
FR-8-9	The web-based Forms will indicate submittal success or failure to users after they submit the information
FR-8-10	The web-based Forms will automatically produce an internal timestamp when information is submitted by an external user
FR 8-11	All decision-making action buttons such as ‘Approve’, ‘Reject’, ‘Cancel’ will include a subsequent ‘Are you sure?’ modal
FR 8-12	Where applicable, the application will adhere to the Americans with Disabilities Act (ADA) & Section 508 recommendations.

5.3 Major Functional Processes

Based on the detailed requirements provided for the major functional portions of the CDAP process, future-state process maps were created to depict the simplified steps and show major portions of the application where automation and validations would be provided. These process maps and descriptions are provided in the next sections and include (1) Vendor Registration, (2) Vendor Credentialing, and (3) Post-Award Documentation Submission & Review.

5.3.1 Vendor Registration

In the future-state process, the vendor will be able to complete a company registration form to obtain their initial company account and administrative user account. The company registration process includes field-level validation to ensure data is being entered in the proper format and error handling ensures that all required fields are completed. The process also includes an email-verification function to ensure that DLA does not get spammed by fake data or suffers other forms of spam-based attacks. As shown in the diagram below, portions of the DLA review are automated, including the email verification and provisioning of the company and user accounts. When an application is successfully submitted, an email notification will be sent to DLA reviewers. Similarly, once an approval/rejection decision has been made, the applicant is also notified via email.

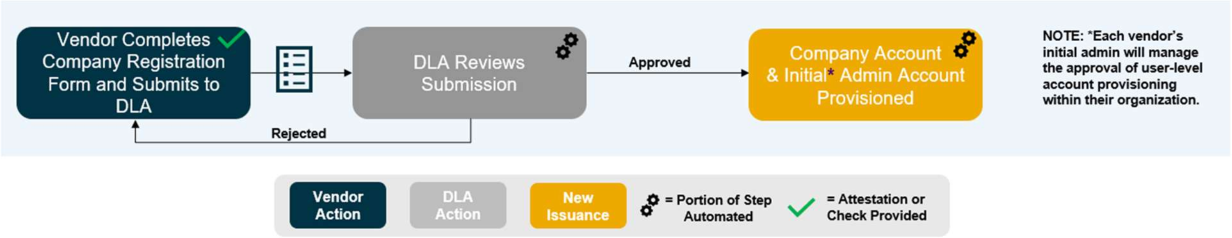
Vendor Completes Company Registration Form and Submits to DLA
DLA Reviews Submission
Approved
Rejected
Company Account & Initial* Admin Account Provisioned
NOTE: *Each vendor's initial admin will manage the approval of user-level account provisioning within their organization.
Vendor Action
DLA Action
New Issuance
= Portion of Step Automated
= Attestation or Check Provided

Figure 15. Process Map of Vendor Registration

As noted in the diagram, the initial issuance of all company accounts will automatically generate a single administrative account for that company based on the user email and information provided in the registration form. This administrator will be responsible for reviewing and approving all subsequent user account applications for this company. Note that company accounts are created for unique Commercial and Government Entity Codes (CAGE) as entered on the registration form.

One unique aspect of the vendor registration process is the creation of a Decentralized Identifier (DID) to associate the company (and CAGE) with a unique identity on the blockchain. This process will be described in more detail during the prototype application feature discussion, but instructions for registering the DID are included in the 'approval' email to successful vendor applicants.

5.3.2 Vendor Credentialing

Perhaps the most simplified step in the future state process is the issuance of vendor quality credentials. This process will be streamlined by utilizing on-file demographic information for the vendor and only requiring them to choose which qualification credential they are applying for and attach documentation. Form validation will ensure that documentation is attached before routing the application to CDAP reviewers as shown in the diagram below.

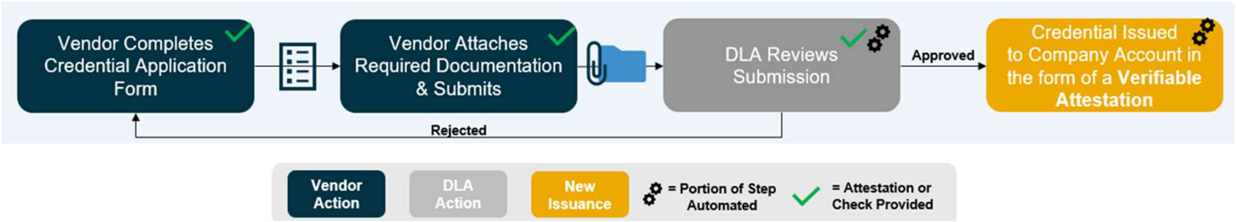
Vendor Completes Credential Application Form
Vendor Attaches Required Documentation & Submits
DLA Reviews Submission
Approved
Credential Issued to Company Account in the form of a Verifiable Attestation
Rejected
Vendor Action
DLA Action
New Issuance
= Portion of Step Automated
= Attestation or Check Provided

Figure 16. Process Map of Vendor Credentialing

DLA reviewers will be notified via email alert when new applications have been received and can focus on reviewing the attached documentation as all other information will be pulled from the system of record as pre-populated for the vendor applicant. Upon reviewing the documentation and making a determination, notification of the approval or rejection is sent to the applicant. If the application for a credential is approved, the qualification credential will be written to the blockchain and provide an easily verifiable attachment for future post-award submissions.

5.3.3 Post-Award Documentation Submission & Review

The most repetitive portion of the CDAP process, but also one of the most important, includes the submission of qualification credentials, Form 918 information, and 'additional documentation 'for CDAP review. This process is triggered whenever a vendor receives an award within Federal Supply Class (FSC) 5962 – the class of components reviewed by CDAP. When vendors win an award, they have a period of time to send the aforementioned data to the CDAP team for review and must receive 'Authorization to Ship' before they can send in the items and complete the process. In the future-state application, vendors will have workload files appear for each new award and will be notified via email whenever an award is made to them. Vendors will then see the awards within their application screens and be able to provide procurement-specific information such as the Manufacturer CAGE, Lot Numbers, Lot Dates, etc. Nearly two-thirds of the previously required Form 918 information (largely demographic and procurement information) will be automatically generated based on stored information on-file.

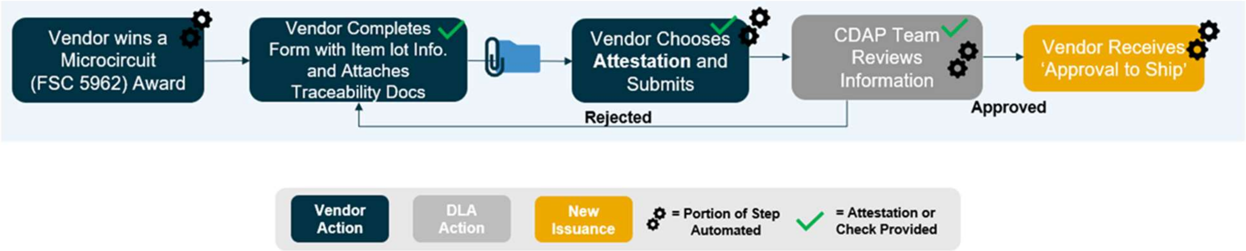
Vendor wins a Microcircuit (FSC 5962) Award
Vendor Completes Form with Item lot Info. and Attaches Traceability Docs
Vendor Chooses Attestation and Submits
CDAP Team Reviews Information
Vendor Receives 'Approval to Ship'
Rejected
Approved
Vendor Action
DLA Action
New Issuance
= Portion of Step Automated
= Attestation or Check Provided

Figure 17. Process Map of Post-Award Documentation Submission & Review

Once a vendor has entered the select fields required to complete Form 918 information, they can select their qualification credential from a drop-down (only approved and issued credentials will appear in this list) and attach any other supporting information via a file attachment function. This information will be securely sent to DLA, and vendors will be required to digitally sign (via a virtual pen cursor) or otherwise confirm the accuracy of their submitted information. Upon receipt of new documentation the CDAP user will be notified via email and can log in to review the information. The vendor credential will be verifiable, via a real-time verification button, which will be explained in detail within the application features section. Upon approval or rejection of individual documents, or when an authorization decision has been made, the vendor will similarly be notified.

6.0 Platform Evaluation & Selection

After completion of the recommended requirements and high-level process design, the technical design and development was ready to begin. However, blockchain technology can be developed with a number of software packages, so a selection process was required. Since the number of software packages and platforms is always growing and more platforms exist than could reasonably be evaluated thoroughly, an approach for slimming down the evaluation pool was required. The following steps were taken to accomplish this process:

1. **Determine Top Platform Selection Criteria** – Custom criteria were developed by thinking about what aspects of a blockchain software platform DLA, and this project, might require. Ultimately, ten criteria were selected to use for evaluation.
2. **Rank the Criteria in Order of Importance –** Using the 'lifeboat method', the criteria were ranked in order of importance. This process involves choosing the most important criterion based on the assumption that you would only be able to keep one for the evaluation, then only two, three, etc. until they are listed in ranked order.
3. **Use the Top Criteria to Down-Select Appropriate Platforms –** Since a large number of software platforms exist, it would not be feasible to reasonably research and evaluate all platforms. Instead, the most important criteria were used to set the standard for which platforms would remain and be evaluated. In other words, if the top criteria were not met, the platform was automatically removed from contention.
4. **Evaluate and Rate Remaining Platforms –** Once the number of platforms was reduced to a reasonable number, detailed research was conducted to provide a relative rating (from zero to four) for each platform within each of the top ten criteria. These ratings were weighted based on their importance and aggregate scores were added up for each platform.
5. **Select Top Platform(s) Based on Stakeholder Input –** The final weighted rankings of available platforms indicated which packages would be the best fit for the criteria developed for this project. Key considerations were explained to the Trusted Working Group and their input was collected in determining the best platform for development.

6.1 Evaluation Criteria

As described above, the top ten criteria for this project were developed based on considerations for both DLA and the overall project. Note that criterion (1) Private Blockchain and (2) Use Case Applicability were used to down-select software platforms to the final six choices. Many of the excluded software packages were either focused on cryptocurrency or other use case areas such as finance. The top ten criteria are listed in order of importance (and weighting):

1. **Private Blockchain** – ***Does the platform support the development of private blockchains?*** The business rules, performance, scalability, regulatory rules, and policies do not align with the public and permissionless model because of data privacy. Therefore, the blockchain platform selected must be private. Due to the nature of DLA's work and information, we will absolutely need to develop a private blockchain and not have all information sitting publicly.
2. **Use Case Applicability** – ***Does the platform support identity/credentialing and/or track and trace?*** The CDAP processes may continue to evolve over time to include both identity and track and trace capabilities, so an ideal platform should be prepared account for both

types of use cases. However, since this prototype is focused on developing an identity-based solution, that use case is most important.

3. **Interoperability** – ***Does the platform have interoperability capabilities?*** The platform needs to have the capability to be interoperable with enterprise systems and perhaps other blockchains to provide value in a production environment. This criterion ensures that the chosen platform will have interoperable capabilities should it proceed to a pilot or production scale implementation.
4. **Production-Readiness –** ***Does the platform support enterprise or other production-ready applications?*** There is the potential that the prototype could eventually move into a live production phase pending positive results. This criterion considers whether the platforms can support enterprise and production-ready applications.
5. **Security Considerations –** ***Does the platform provide a technical approach that ensures proper data control and security?*** Most blockchain platforms utilize the same cryptographic approaches, but the technical design may expose data and/or access in undesirable ways. Given DLA's focus on cybersecurity, these considerations need to be highly considered.
6. **Platform Maturity** – ***Has the platform been well established, has a mature consortium, and/or mature development tools available?*** New platforms are always emerging, but an ideal platform will have strong industry/development support and multiple examples of use cases built around it. Similarly, it is important to understand whether a platform is mature enough to include development tools and/or technical support.
7. **Development Complexity –** ***Is a very specialized skillset and/or proprietary tools/access required to develop on this platform?*** An ideal platform should utilize well-known or common development tools, allowing for intuitive handoff between any number of current or future developers. Platforms with very esoteric or complex development types would not be ideal for this type of use case.
8. **Smart Contract Enabled? –** ***Does the platform support Smart Contracts and automated Execution Logic?*** There are several exchanges in the current process that are standard, straightforward, and repeatable. Using smart contracts will allow DLA to automate aspects of the process and provide real efficiency and value to the client and their external partners. Smart contracts allow an entity to have a shared process across multiple organizations and ensure that the transitions happen according to a predefined set of rules.
9. **Future Flexibility –** ***Does the platform provide some flexibility for modifying technical/functional features or adding in capabilities in the future?*** This use case may grow to incorporate a larger group of stakeholders and additional functionality in future phases. An ideal platform will allow flexibility for future development and not be limited by the platform.
10. **Governance & Consensus –** ***Does the platform have flexibility with how Governance and Consensus work, to fit to our use case?*** As described in the Governance section, the approach for governing this application may develop over time, starting as a fully-owned DLA capability. If future government partners are onboarded, some flexibility may be required in the way governance is handled from a technical perspective, including the types

of consensus mechanisms available – the means by which major stakeholder groups agree to add valid transactions to the blockchain ledger.

6.2 Platform Evaluations

Based on the down-selection described in Step 3 of the overall process, six final blockchain development platforms were researched and rated according to the top ten criteria explained in the previous section.

These final six platforms include (1) Hyperledger Indy, (2) Hyperledger Fabric, (3) Corda, (4) Hyperledger Sawtooth, (5) Quorum, and (6) Guardtime. Based on available research results, each platform was rated on a scale from zero to four and represented by 'Harvey Balls' as shown in the evaluation legend below.

Table 11. Harvey Ball Rating Legend for Platform Evaluation

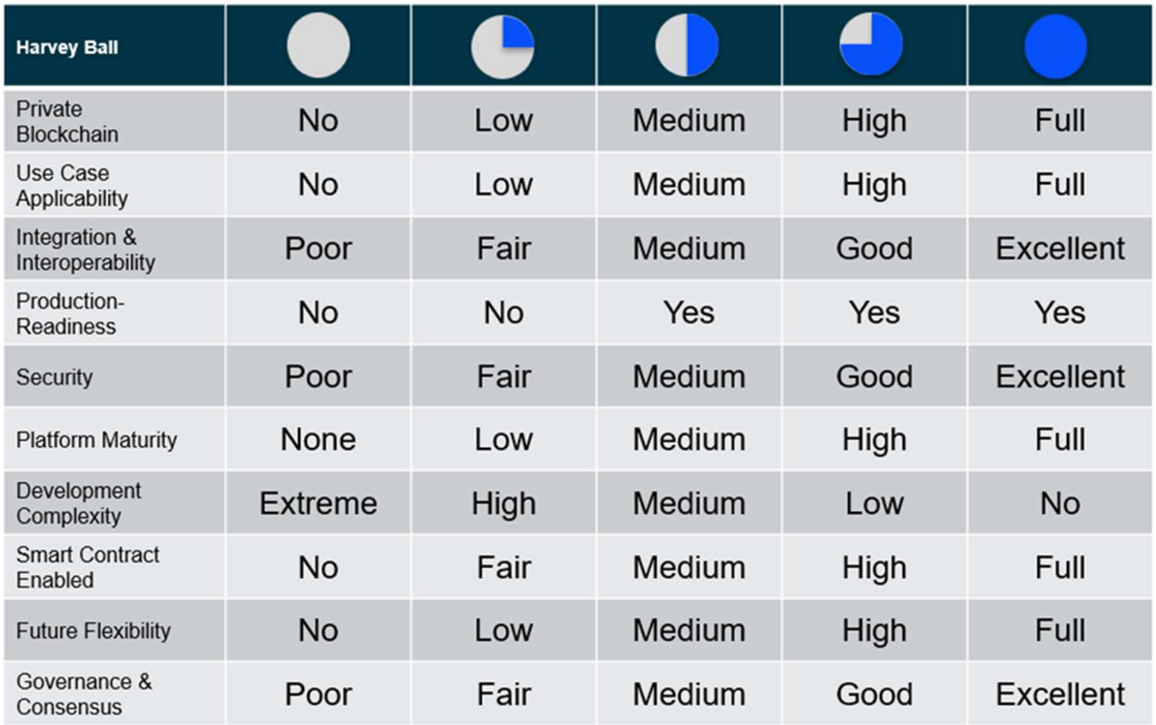

Harvey Ball	○	◔	◑	◕	●
Private Blockchain	No	Low	Medium	High	Full
Use Case Applicability	No	Low	Medium	High	Full
Integration & Interoperability	Poor	Fair	Medium	Good	Excellent
Production-Readiness	No	No	Yes	Yes	Yes
Security	Poor	Fair	Medium	Good	Excellent
Platform Maturity	None	Low	Medium	High	Full
Development Complexity	Extreme	High	Medium	Low	No
Smart Contract Enabled	No	Fair	Medium	High	Full
Future Flexibility	No	Low	Medium	High	Full
Governance & Consensus	Poor	Fair	Medium	Good	Excellent

Note: Half of the final blockchain platforms evaluated and rated are projects developed by the Hyperledger Consortium – these are the three platforms that begin with the prefix 'Hyperledger.' This is not surprising, because the Hyperledger consortium is a top collaborative effort within the blockchain space hosted by the Linux Foundation, created to advance cross-industry blockchain technologies and development. Major blockchain-focused companies including Accenture, IBM, Intel, and SAP (among others) have collaborated to advance several of the platforms on an open source basis, moving forward packages primarily relevant to specific private business blockchain use cases.

Detailed evaluation ratings and their corresponding explanations are shown in tables on the following pages. More comprehensive descriptions of each platform and their ratings are included in the Appendices.

Table 12. Platform Evaluation Results for Hyperledger Indy

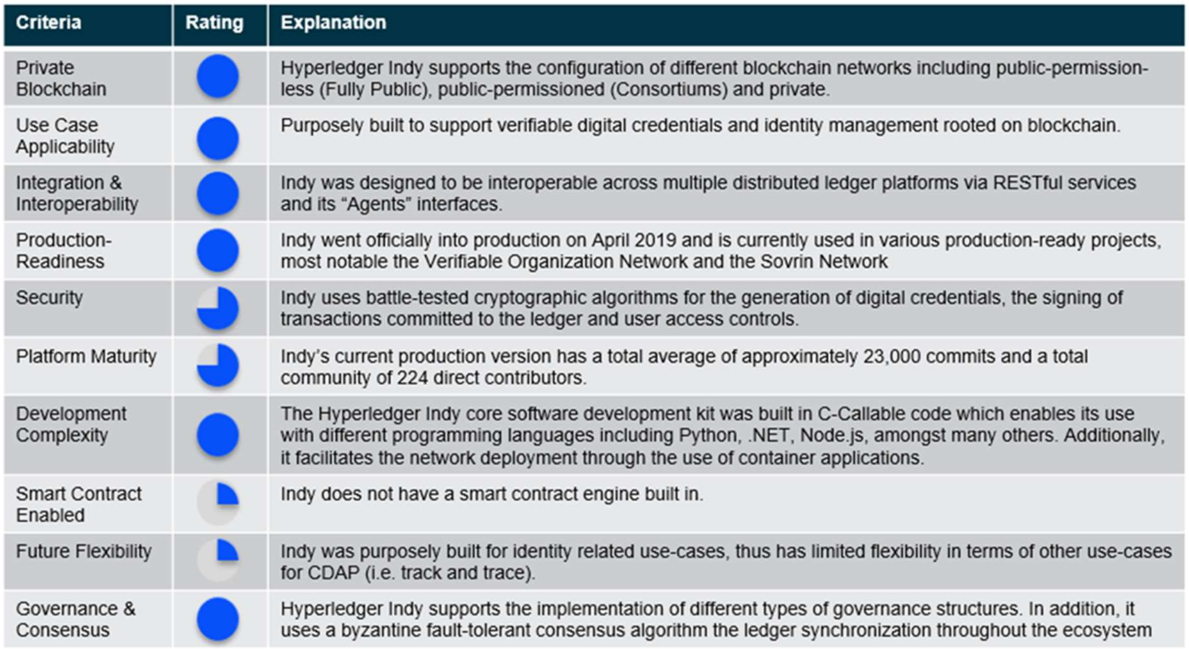

Criteria	Rating	Explanation
Private Blockchain	●	Hyperledger Indy supports the configuration of different blockchain networks including public-permission-less (Fully Public), public-permissioned (Consortiums) and private.
Use Case Applicability	●	Purposely built to support verifiable digital credentials and identity management rooted on blockchain.
Integration & Interoperability	●	Indy was designed to be interoperable across multiple distributed ledger platforms via RESTful services and its "Agents" interfaces.
Production-Readiness	●	Indy went officially into production on April 2019 and is currently used in various production-ready projects, most notable the Verifiable Organization Network and the Sovrin Network
Security	◕	Indy uses battle-tested cryptographic algorithms for the generation of digital credentials, the signing of transactions committed to the ledger and user access controls.
Platform Maturity	◕	Indy's current production version has a total average of approximately 23,000 commits and a total community of 224 direct contributors.
Development Complexity	●	The Hyperledger Indy core software development kit was built in C-Callable code which enables its use with different programming languages including Python, .NET, Node.js, amongst many others. Additionally, it facilitates the network deployment through the use of container applications.
Smart Contract Enabled	◔	Indy does not have a smart contract engine built in.
Future Flexibility	◔	Indy was purposely built for identity related use-cases, thus has limited flexibility in terms of other use-cases for CDAP (i.e. track and trace).
Governance & Consensus	●	Hyperledger Indy supports the implementation of different types of governance structures. In addition, it uses a byzantine fault-tolerant consensus algorithm the ledger synchronization throughout the ecosystem

Table 13. Platform Evaluation Results for Hyperledger Fabric

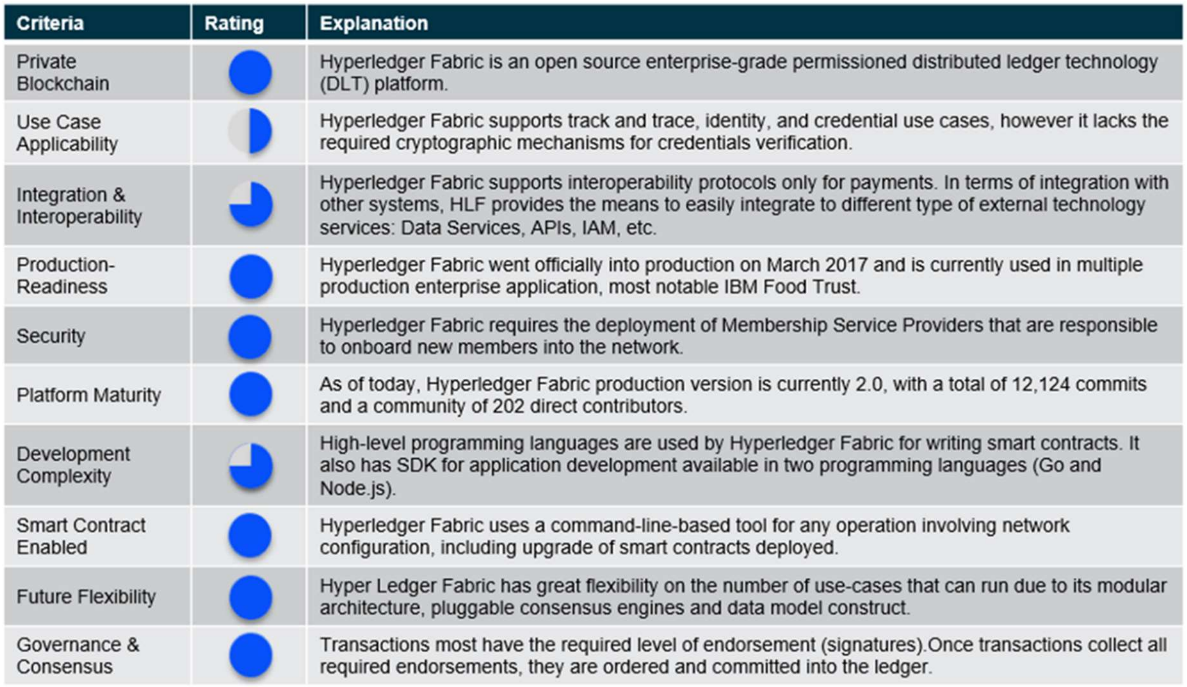

Criteria	Rating	Explanation
Private Blockchain	●	Hyperledger Fabric is an open source enterprise-grade permissioned distributed ledger technology (DLT) platform.
Use Case Applicability	◑	Hyperledger Fabric supports track and trace, identity, and credential use cases, however it lacks the required cryptographic mechanisms for credentials verification.
Integration & Interoperability	◕	Hyperledger Fabric supports interoperability protocols only for payments. In terms of integration with other systems, HLF provides the means to easily integrate to different type of external technology services: Data Services, APIs, IAM, etc.
Production-Readiness	●	Hyperledger Fabric went officially into production on March 2017 and is currently used in multiple production enterprise application, most notable IBM Food Trust.
Security	●	Hyperledger Fabric requires the deployment of Membership Service Providers that are responsible to onboard new members into the network.
Platform Maturity	●	As of today, Hyperledger Fabric production version is currently 2.0, with a total of 12,124 commits and a community of 202 direct contributors.
Development Complexity	◕	High-level programming languages are used by Hyperledger Fabric for writing smart contracts. It also has SDK for application development available in two programming languages (Go and Node.js).
Smart Contract Enabled	●	Hyperledger Fabric uses a command-line-based tool for any operation involving network configuration, including upgrade of smart contracts deployed.
Future Flexibility	●	Hyper Ledger Fabric has great flexibility on the number of use-cases that can run due to its modular architecture, pluggable consensus engines and data model construct.
Governance & Consensus	●	Transactions most have the required level of endorsement (signatures).Once transactions collect all required endorsements, they are ordered and committed into the ledger.

Table 14. Platform Evaluation Results for Corda

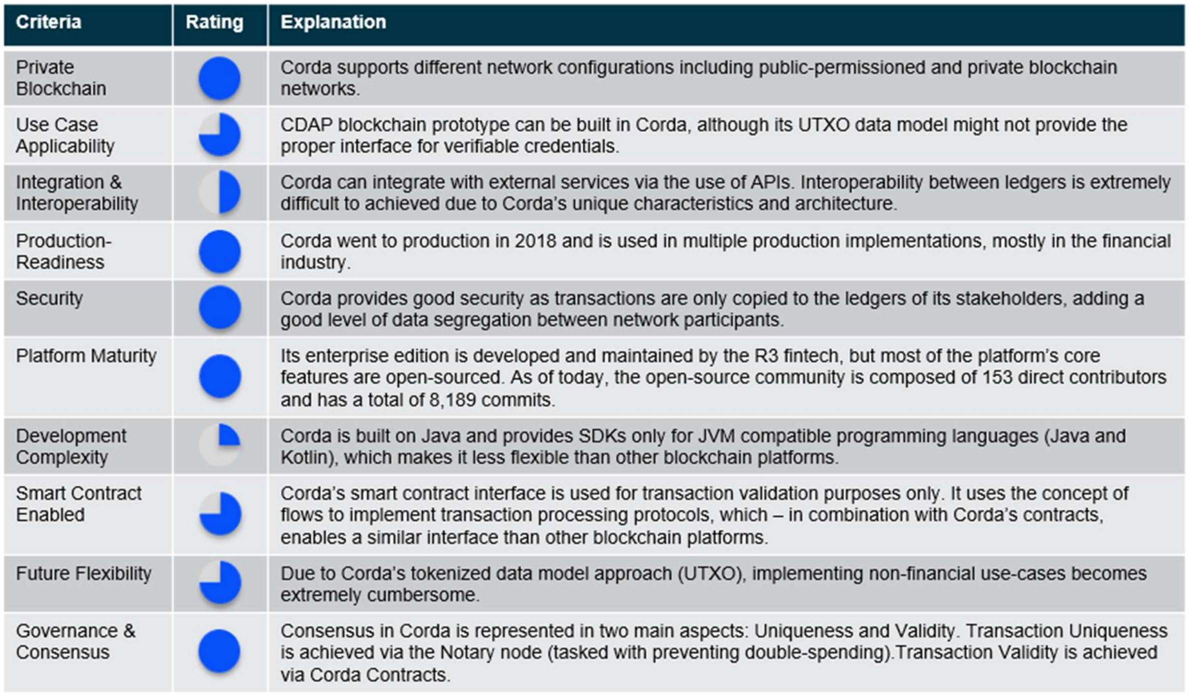

Criteria	Rating	Explanation
Private Blockchain	●	Corda supports different network configurations including public-permissioned and private blockchain networks.
Use Case Applicability	◕	CDAP blockchain prototype can be built in Corda, although its UTXO data model might not provide the proper interface for verifiable credentials.
Integration & Interoperability	◑	Corda can integrate with external services via the use of APIs. Interoperability between ledgers is extremely difficult to achieved due to Corda's unique characteristics and architecture.
Production-Readiness	●	Corda went to production in 2018 and is used in multiple production implementations, mostly in the financial industry.
Security	●	Corda provides good security as transactions are only copied to the ledgers of its stakeholders, adding a good level of data segregation between network participants.
Platform Maturity	●	Its enterprise edition is developed and maintained by the R3 fintech, but most of the platform's core features are open-sourced. As of today, the open-source community is composed of 153 direct contributors and has a total of 8,189 commits.
Development Complexity	◔	Corda is built on Java and provides SDKs only for JVM compatible programming languages (Java and Kotlin), which makes it less flexible than other blockchain platforms.
Smart Contract Enabled	◕	Corda's smart contract interface is used for transaction validation purposes only. It uses the concept of flows to implement transaction processing protocols, which – in combination with Corda's contracts, enables a similar interface than other blockchain platforms.
Future Flexibility	◕	Due to Corda's tokenized data model approach (UTXO), implementing non-financial use-cases becomes extremely cumbersome.
Governance & Consensus	●	Consensus in Corda is represented in two main aspects: Uniqueness and Validity. Transaction Uniqueness is achieved via the Notary node (tasked with preventing double-spending).Transaction Validity is achieved via Corda Contracts.

Table 15. Platform Evaluation Results for Hyperledger Sawtooth

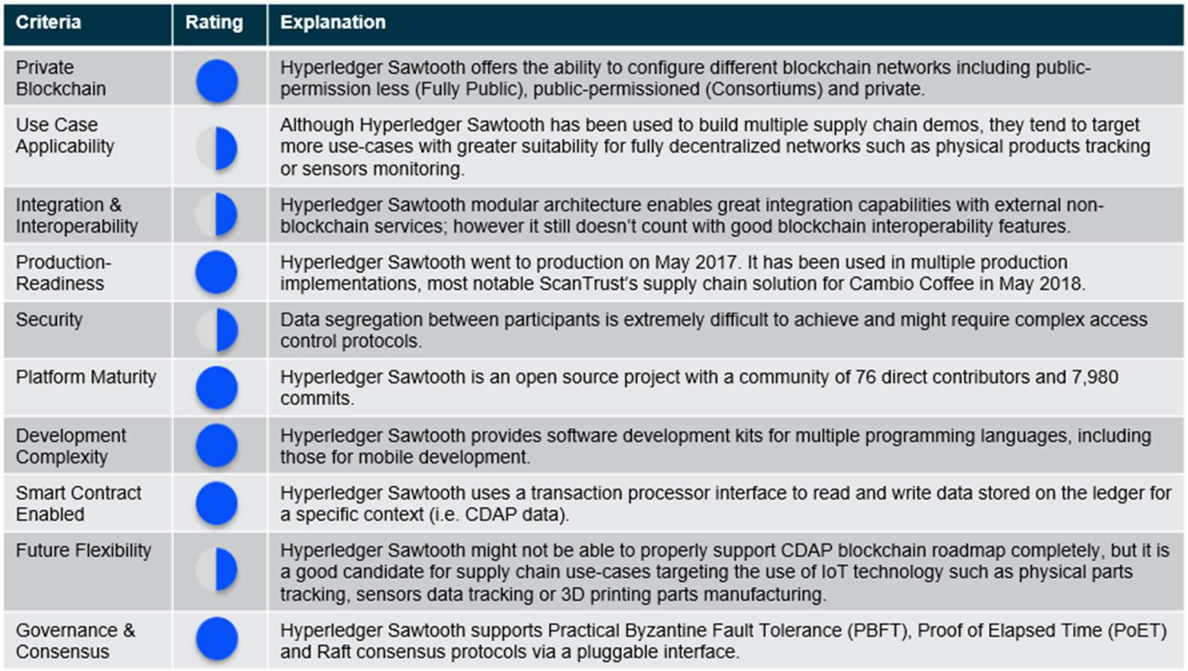

Criteria	Rating	Explanation
Private Blockchain	●	Hyperledger Sawtooth offers the ability to configure different blockchain networks including public-permission less (Fully Public), public-permissioned (Consortiums) and private.
Use Case Applicability	◑	Although Hyperledger Sawtooth has been used to build multiple supply chain demos, they tend to target more use-cases with greater suitability for fully decentralized networks such as physical products tracking or sensors monitoring.
Integration & Interoperability	◑	Hyperledger Sawtooth modular architecture enables great integration capabilities with external non-blockchain services; however it still doesn't count with good blockchain interoperability features.
Production-Readiness	●	Hyperledger Sawtooth went to production on May 2017. It has been used in multiple production implementations, most notable ScanTrust's supply chain solution for Cambio Coffee in May 2018.
Security	◑	Data segregation between participants is extremely difficult to achieve and might require complex access control protocols.
Platform Maturity	●	Hyperledger Sawtooth is an open source project with a community of 76 direct contributors and 7,980 commits.
Development Complexity	●	Hyperledger Sawtooth provides software development kits for multiple programming languages, including those for mobile development.
Smart Contract Enabled	●	Hyperledger Sawtooth uses a transaction processor interface to read and write data stored on the ledger for a specific context (i.e. CDAP data).
Future Flexibility	◑	Hyperledger Sawtooth might not be able to properly support CDAP blockchain roadmap completely, but it is a good candidate for supply chain use-cases targeting the use of IoT technology such as physical parts tracking, sensors data tracking or 3D printing parts manufacturing.
Governance & Consensus	●	Hyperledger Sawtooth supports Practical Byzantine Fault Tolerance (PBFT), Proof of Elapsed Time (PoET) and Raft consensus protocols via a pluggable interface.

Table 16. Platform Evaluation Results for Quorum

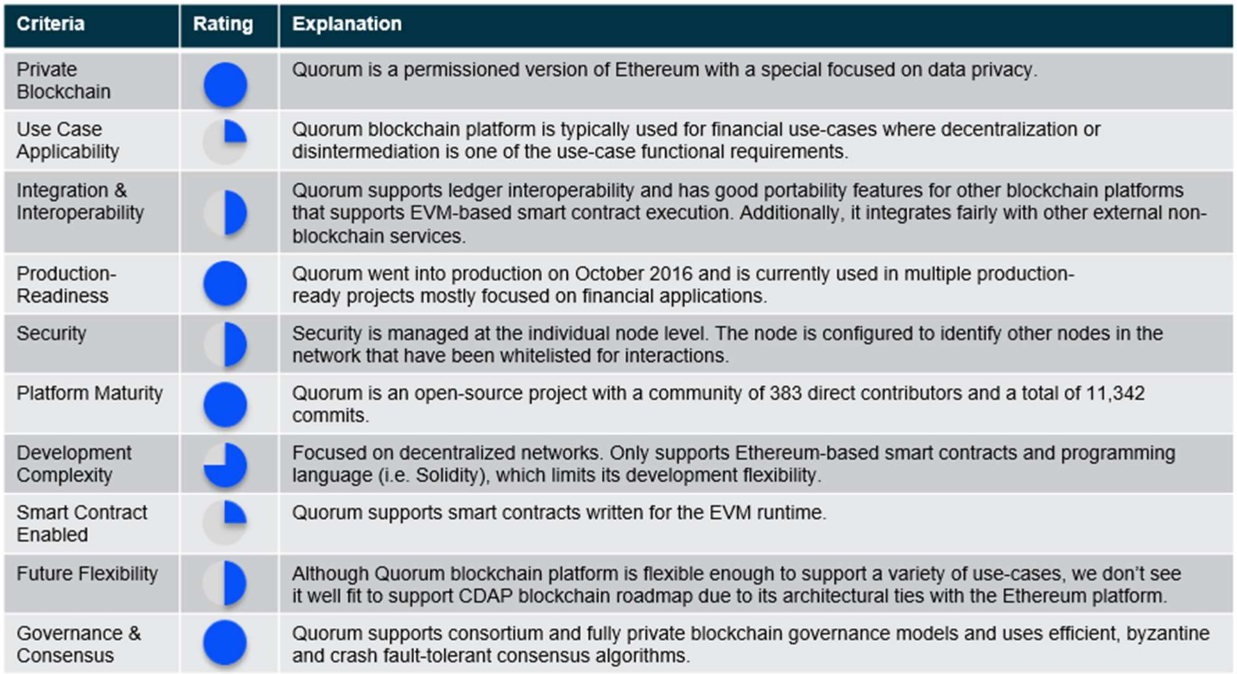

Criteria	Rating	Explanation
Private Blockchain	●	Quorum is a permissioned version of Ethereum with a special focused on data privacy.
Use Case Applicability	◔	Quorum blockchain platform is typically used for financial use-cases where decentralization or disintermediation is one of the use-case functional requirements.
Integration & Interoperability	◑	Quorum supports ledger interoperability and has good portability features for other blockchain platforms that supports EVM-based smart contract execution. Additionally, it integrates fairly with other external non-blockchain services.
Production-Readiness	●	Quorum went into production on October 2016 and is currently used in multiple production-ready projects mostly focused on financial applications.
Security	◑	Security is managed at the individual node level. The node is configured to identify other nodes in the network that have been whitelisted for interactions.
Platform Maturity	●	Quorum is an open-source project with a community of 383 direct contributors and a total of 11,342 commits.
Development Complexity	◕	Focused on decentralized networks. Only supports Ethereum-based smart contracts and programming language (i.e. Solidity), which limits its development flexibility.
Smart Contract Enabled	◔	Quorum supports smart contracts written for the EVM runtime.
Future Flexibility	◑	Although Quorum blockchain platform is flexible enough to support a variety of use-cases, we don't see it well fit to support CDAP blockchain roadmap due to its architectural ties with the Ethereum platform.
Governance & Consensus	●	Quorum supports consortium and fully private blockchain governance models and uses efficient, byzantine and crash fault-tolerant consensus algorithms.

Table 17. Platform Evaluation for Guardtime

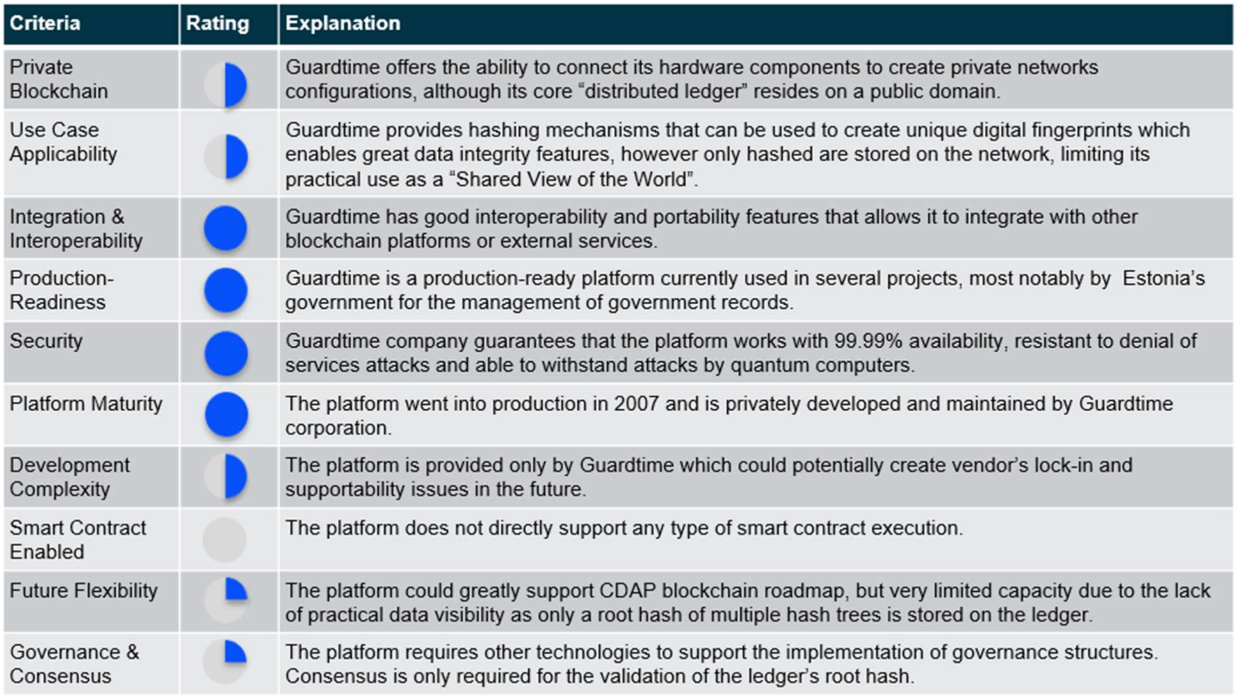

Criteria	Rating	Explanation
Private Blockchain	◑	Guardtime offers the ability to connect its hardware components to create private networks configurations, although its core "distributed ledger" resides on a public domain.
Use Case Applicability	◑	Guardtime provides hashing mechanisms that can be used to create unique digital fingerprints which enables great data integrity features, however only hashed are stored on the network, limiting its practical use as a "Shared View of the World".
Integration & Interoperability	●	Guardtime has good interoperability and portability features that allows it to integrate with other blockchain platforms or external services.
Production-Readiness	●	Guardtime is a production-ready platform currently used in several projects, most notably by Estonia's government for the management of government records.
Security	●	Guardtime company guarantees that the platform works with 99.99% availability, resistant to denial of services attacks and able to withstand attacks by quantum computers.
Platform Maturity	●	The platform went into production in 2007 and is privately developed and maintained by Guardtime corporation.
Development Complexity	◑	The platform is provided only by Guardtime which could potentially create vendor's lock-in and supportability issues in the future.
Smart Contract Enabled	○	The platform does not directly support any type of smart contract execution.
Future Flexibility	◔	The platform could greatly support CDAP blockchain roadmap, but very limited capacity due to the lack of practical data visibility as only a root hash of multiple hash trees is stored on the ledger.
Governance & Consensus	◔	The platform requires other technologies to support the implementation of governance structures. Consensus is only required for the validation of the ledger's root hash.

6.3 Platform Evaluation Comparison

The results of the weighted platform evaluation are shown below, where criteria are listed from left to right on the x-axis in decreasing order of importance and platforms are listed from top to bottom on the y-axis in decreasing order of overall weighted score. In general, the top two platforms, Hyperledger Indy and Hyperledger Fabric scored the best and were presented to the Trusted Working Group (TWG) as the primary platforms to consider for development.

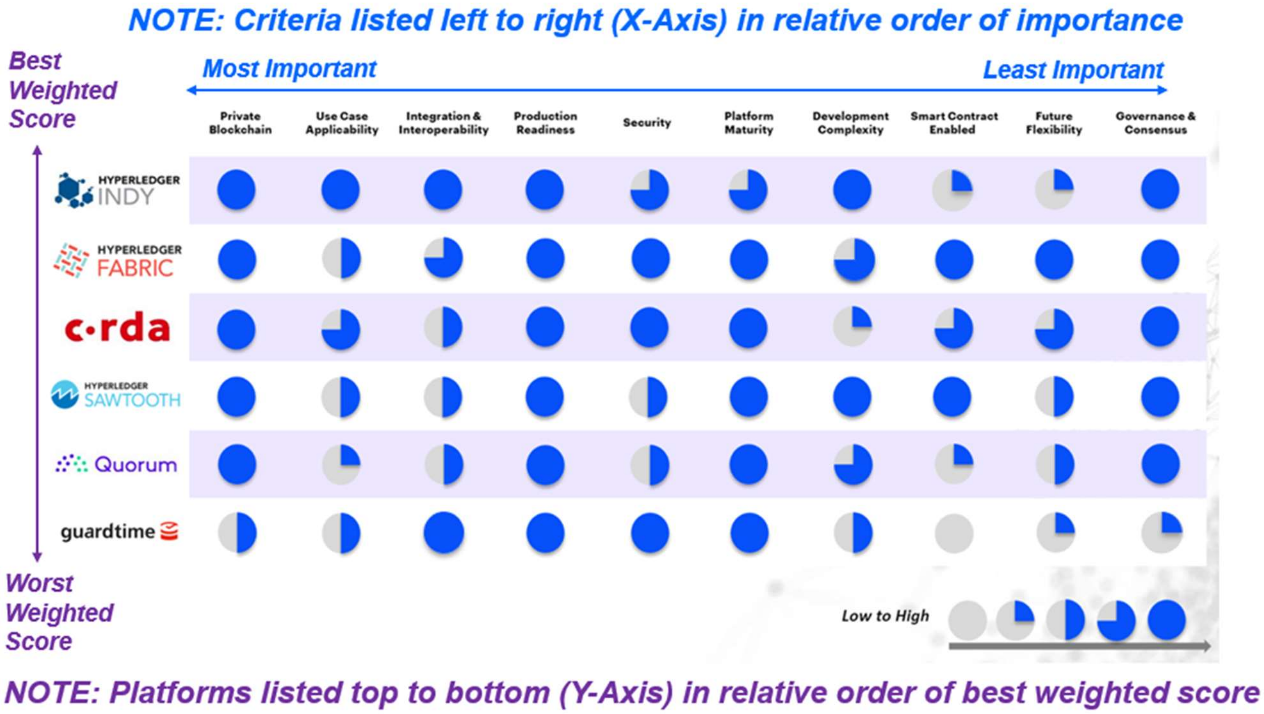

NOTE: Criteria listed left to right (X-Axis) in relative order of importance

Best Weighted Score

Most Important → Least Important

	Private Blockchain	Use Case Applicability	Integration & Interoperability	Production Readiness	Security	Platform Maturity	Development Complexity	Smart Contract Enabled	Future Flexibility	Governance & Consensus
HYPERLEDGER INDY	●	●	●	●	◕	◕	●	◔	◔	●
HYPERLEDGER FABRIC	●	◑	◕	●	●	●	◕	●	●	●
c•rda	●	◕	◑	●	●	●	◔	◕	◕	●
HYPERLEDGER SAWTOOTH	●	◑	◑	●	◑	●	●	●	◑	●
Quorum	●	◔	◑	●	◑	●	◕	◔	◑	●
guardtime	◑	◑	●	●	●	●	◑	○	◔	◔

Worst Weighted Score

Low to High ○ ◔ ◑ ◕ ●

NOTE: Platforms listed top to bottom (Y-Axis) in relative order of best weighted score

Figure 18. Results of Weighted Platform Evaluation

Hyperledger Indy scored well because it is an identity-focused platform developed by the Hyperledger Consortium, so it is extremely applicable to our use case and was further differentiated by scoring very well in areas such as integration, production-readiness, and development complexity. Indy did not score well in smart contracts due to the lack of a pre-existing smart contract engine, however this is not a deal-breaker since they are not required for this use case.

Hyperledger Fabric is a more general-purpose blockchain platform, meant to serve a variety of use cases and provide broad flexibility. For these reasons, it scores well in areas that other platforms did not including integration, security, development complexity, smart contracts, and flexibility. Overall Hyperledger Fabric can be used for a variety of applications but lacks some of the pre-built technical structure and features that make a purpose-built platform like Indy desirable.

6.4 Platform Recommendation

In reviewing the results with the Trusted Working Group, it was determined that ***Hyperledger Indy*** would be used as our software development platform. Hyperledger Indy is a purpose-built distributed ledger platform for decentralized identity and the first identity-focused blockchain framework created within the Hyperledger consortium. It features verifiable credentials based on zero-knowledge proof (ZKP) technology, decentralized identifiers (DIDs), a software development kit (SDK) for building applications, and a node infrastructure to manage the distributed ledger.

In addition to utilizing Hyperledger Indy, it was recommended to the working group by the R&D Program Office and subject matter experts that ***Hyperledger Aries*** be considered for use as well. Hyperledger Aries is a project within the Hyperledger consortium that provides a set of libraries and infrastructure components for the development and deployment identity agent services. Although is currently being built to interact only with Hyperledger Indy, eventually it will support other blockchain identity platforms too. DLA could benefit from using Aries and its interoperable features to avoid getting locked-in with one specific blockchain platform. Although Aries is not a standalone platform, it offers integration and infrastructure components for identity-based blockchain solutions including the Aries Cloud Agent (Python) framework. This framework is an open-source toolkit that provides out-of-the-box features for secured agent-to-agent communication, data encryption, credentials exchange, key management, and data storage.

Ultimately, Aries was found to lack a key component of identity credential management, the ability to perform revocation of a credential. This functional component is key to our use case, because one major advantage of storing vendor credentials on-chain is the ability for the issuing agency (in this case DLA) to revoke the credentials instantly if fraud (or other legal justifications) are determined. In a multi-agency application, this component is paramount to preserving supply chain integrity by allowing component agencies to verify credentials at the time of award and review. Hyperledger Aries is a package that is still in constant development, so it should be revisited and utilized in future phases and/or identity solutions once it is more mature and offers revocation capabilities.

7.0 Technical Approach

In the ***Blockchain Feasibility Study*** completed prior to this project, requirements to enhance the post-award CDAP process using blockchain technology were analyzed and determined technically feasible in a prototype application.

To demonstrate this, a prototype was designed to simulate the CDAP post-award operations in a collaborative environment where different participants, internal and external to DLA, are able to share information securely. After validating functional requirements with business stakeholders, non-functional requirements were identified to establish a baseline for the selection of capabilities and software resources in line with DLA and DoD cybersecurity guidelines. A distinctive reference architecture model was used to develop a microservice architecture, focused on architecture principles of Adaptability and Flexibility, Information Security and Privacy, Governance and Compliance, Integrity and Scalability.

To build the prototype, a team of developers and specialists was assembled to work on the different parts of the prototype's architecture. The development phase was divided in seven sprints, each with a duration of two weeks. During each sprint, the team completed key functionality throughout different parts of the architecture and merged updates to the code to a running instance in the cloud, using a Continuous Integration/Continuous Delivery (CI/CD) approach.

In the following sections, a high-level walkthrough of the reference architecture used to design the prototype will be examined and described with all the services that were developed. Next, technical information about the physical components will be introduced including how they were deployed on the Amazon Web Services (AWS) cloud. Lastly, the technical transition plan will be discussed with considerations for the pilot and production phases, as well as lessons learned during the development of this prototype.

7.1 Technical Prototype Design

A blockchain solution is not entirely based on blockchain technology, rather it is instead a collection of capabilities that are implemented to address the problem at hand. A reference architecture model was used to analyze the functional and non-functional requirements and design the solution technical blueprint. Non-functional requirements were produced by the Technology Architect and Technical Lead during the design phase. They serve as design constraints so that the final solution addresses areas of technical importance. Furthermore, a formal framework was used for non-functional requirements that are specific to blockchain. Figure 19 on the next page shows a summary of this framework, with a high-level description of each non-functional domain.

This framework does not distinguish between a prototype or a production solution, instead it considers a system powered by blockchain holistically. Each area was referenced for the prototype to determine whether it applied to this phase. Key areas of operation were included regardless to maximize the production-readiness of the solution. For example, performance, although primarily a concern for production-ready solutions, was not a key driver for this prototype. However, careful consideration was taken to support high availability requirements – a performance indicator – through the use Kubernetes container-orchestration technology for the deployment, scaling and management of application components.

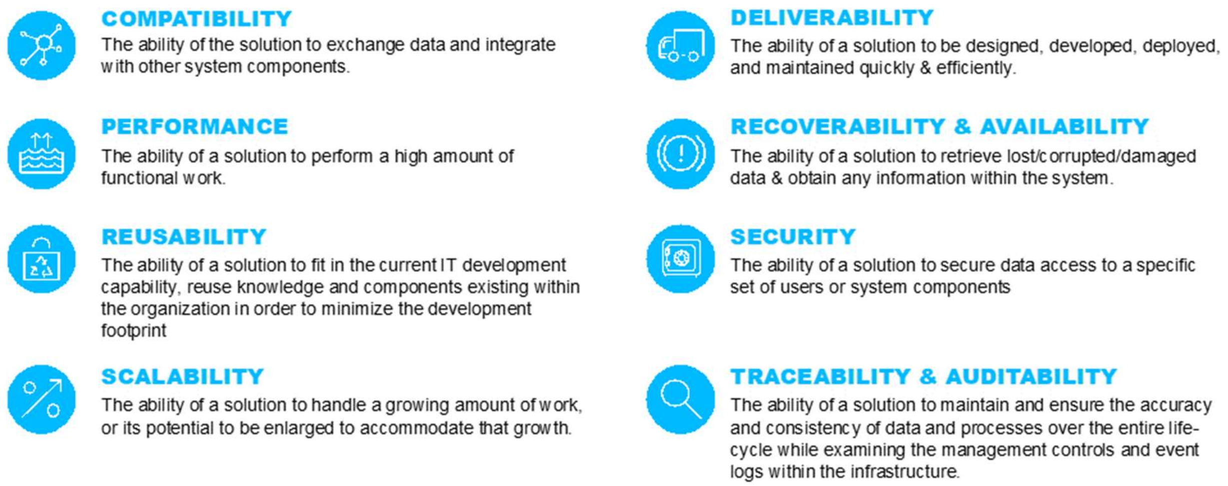

COMPATIBILITY
The ability of the solution to exchange data and integrate with other system components.

PERFORMANCE
The ability of a solution to perform a high amount of functional work.

REUSABILITY
The ability of a solution to fit in the current IT development capability, reuse knowledge and components existing within the organization in order to minimize the development footprint

SCALABILITY
The ability of a solution to handle a growing amount of work, or its potential to be enlarged to accommodate that growth.

DELIVERABILITY
The ability of a solution to be designed, developed, deployed, and maintained quickly & efficiently.

RECOVERABILITY & AVAILABILITY
The ability of a solution to retrieve lost/corrupted/damaged data & obtain any information within the system.

SECURITY
The ability of a solution to secure data access to a specific set of users or system components

TRACEABILITY & AUDITABILITY
The ability of a solution to maintain and ensure the accuracy and consistency of data and processes over the entire life-cycle while examining the management controls and event logs within the infrastructure.

Figure 19. Framework of Non-Functional Requirements for Blockchain Solutions

Figure 20 below shows a list of non-functional requirements that we identified as key requirements for the prototype. Security, traceability, and auditability were major drivers of the prototype design. The final solution must have capabilities to manage access control granularly across different types of users and organizations, as well as securing all application programming interface endpoints with enterprise-grade authentication protocols. Additionally, privacy was paramount for our design with specific guidelines to prevent storing sensitive or confidential information on the blockchain ledger.

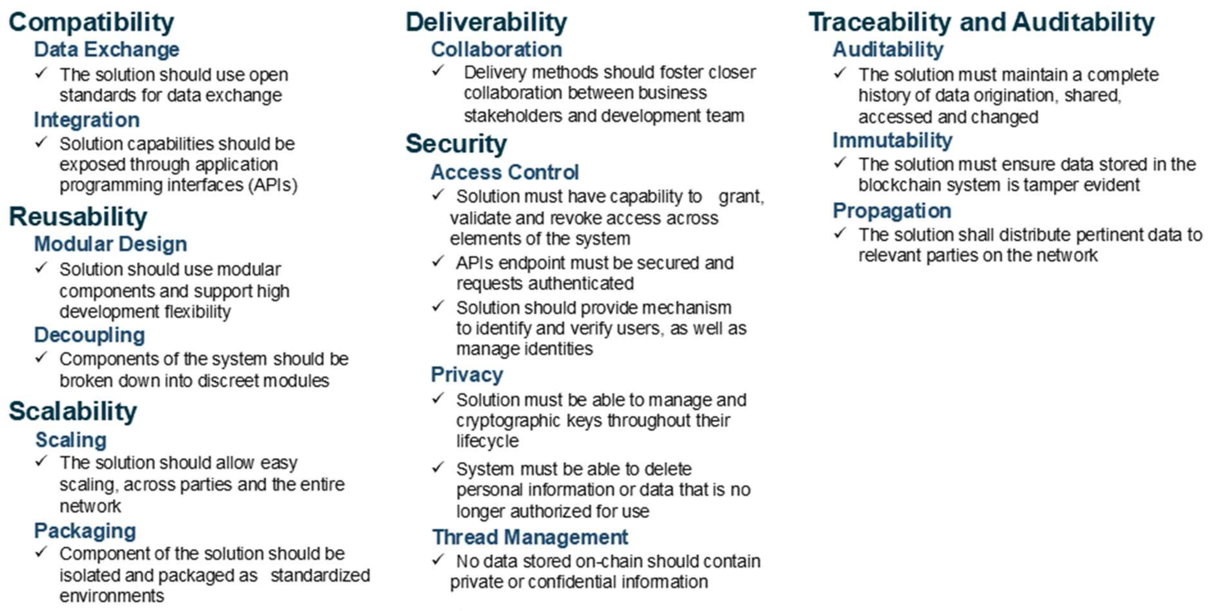

Compatibility
Data Exchange
✓ The solution should use open standards for data exchange
Integration
✓ Solution capabilities should be exposed through application programming interfaces (APIs)

Reusability
Modular Design
✓ Solution should use modular components and support high development flexibility
Decoupling
✓ Components of the system should be broken down into discreet modules

Scalability
Scaling
✓ The solution should allow easy scaling, across parties and the entire network
Packaging
✓ Component of the solution should be isolated and packaged as standardized environments

Deliverability
Collaboration
✓ Delivery methods should foster closer collaboration between business stakeholders and development team

Security
Access Control
✓ Solution must have capability to grant, validate and revoke access across elements of the system
✓ APIs endpoint must be secured and requests authenticated
✓ Solution should provide mechanism to identify and verify users, as well as manage identities
Privacy
✓ Solution must be able to manage and cryptographic keys throughout their lifecycle
✓ System must be able to delete personal information or data that is no longer authorized for use
Thread Management
✓ No data stored on-chain should contain private or confidential information

Traceability and Auditability
Auditability
✓ The solution must maintain a complete history of data origination, shared, accessed and changed
Immutability
✓ The solution must ensure data stored in the blockchain system is tamper evident
Propagation
✓ The solution shall distribute pertinent data to relevant parties on the network

Figure 20. Non-Functional Requirements Identified for the Prototype

The solution must also provide means to perform audits, thus the need to maintain a complete history of data and how it was originated, shared, accessed and updated. Another major requirement was building a solution with a high degree of modularity and decoupling, which will allow DLA to easily maintain and even repurpose components of the application, if needed for production readiness and compliance, while also ensuring that impacts to other areas of the system were minimized and controlled.

Compatibility, specifically in the use of data standards in compliance with DLA guidelines, was also a key requirement for the prototype's design. Furthermore, the use of data standards will not only apply to those specific to DLA, but also standards defined by external entities, such as the World Wide Web Consortium (W3C) data standards for Verifiable Credentials and Decentralized Identifiers, used in the creation and management of digital credentials rooted in blockchain. Leveraging formal data standard models also increases the reciprocity of the information shared within the ecosystem, as well as the level of interoperability with enterprise systems, either internal or external. The ability to integrate the solution to enterprise system was of great importance, specifically to advance this project to a pilot phase.

As previously stated, we use the reference architecture illustrated in Figure 21 to organize capabilities by service areas. These areas include Presentation, Data, Integration, DevOps, Security, Infrastructure, and the Distributed Ledger Technology (DLT) or the blockchain platform. Within each service area we identified the different capabilities required to address the solution's functional and non-functional requirements. Note that the reference architecture also differentiates development and operational services versus runtime services. Development and operational services are those are those used either during the development of the solution or established administratively via predefined protocols and rules. Runtime services are those that include capabilities that are required in order for the solution to work as expected.

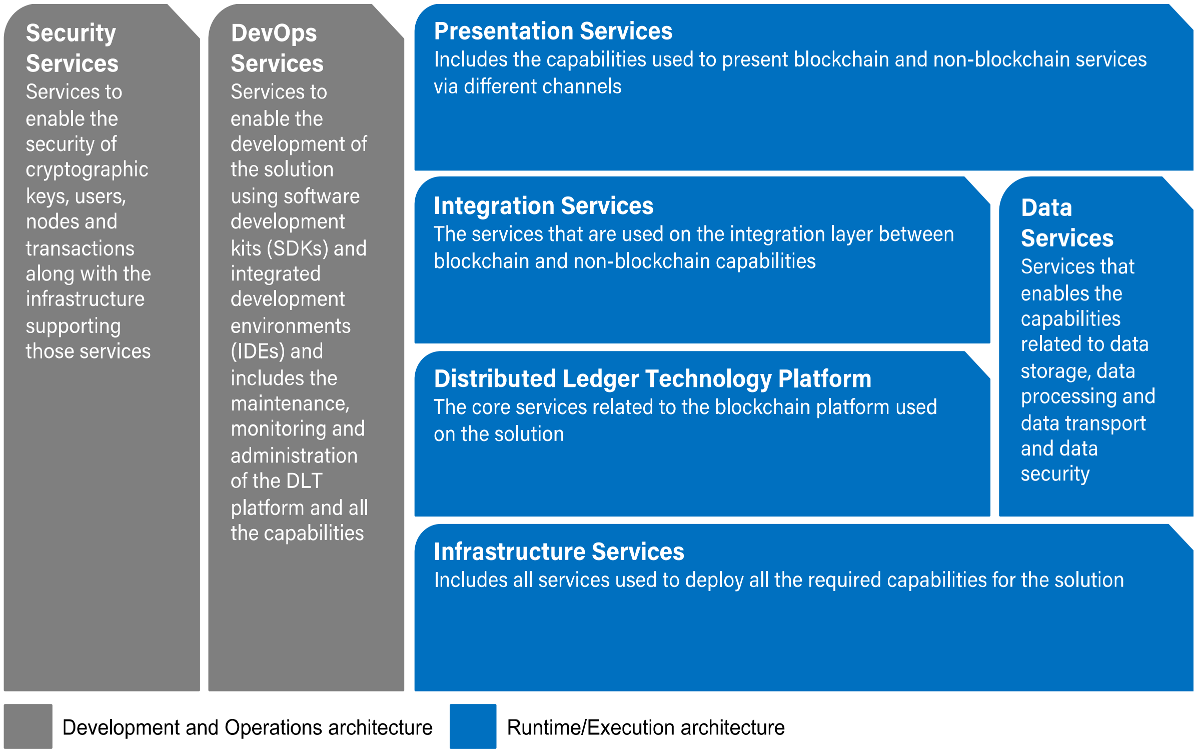
Security Services
Services to enable the security of cryptographic keys, users, nodes and transactions along with the infrastructure supporting those services
DevOps Services
Services to enable the development of the solution using software development kits (SDKs) and integrated development environments (IDEs) and includes the maintenance, monitoring and administration of the DLT platform and all the capabilities
Presentation Services
Includes the capabilities used to present blockchain and non-blockchain services via different channels
Integration Services
The services that are used on the integration layer between blockchain and non-blockchain capabilities
Data Services
Services that enables the capabilities related to data storage, data processing and data transport and data security
Distributed Ledger Technology Platform
The core services related to the blockchain platform used on the solution
Infrastructure Services
Includes all services used to deploy all the required capabilities for the solution
Development and Operations architecture
Runtime/Execution architecture

Figure 21. Reference Architecture Model for Blockchain Solutions

To assist in the selection of capabilities we established architecture principles that, in combination with non-functional requirements, provided a blueprint for technical decisions making in the development of the prototype's architecture.

An architecture principle is a statement of belief, approach or intent which directs the formulation of the architecture. These principles serve as guidelines for the construction of the architecture. They are aligned with the non-functional requirements and used to justify the decisions made about the patterns, components, and capabilities in the architecture, ensuring that the architecture defined is consistent. The figure below shows a list of high-level architecture principles often considered for blockchain-based solutions along with a summarized description

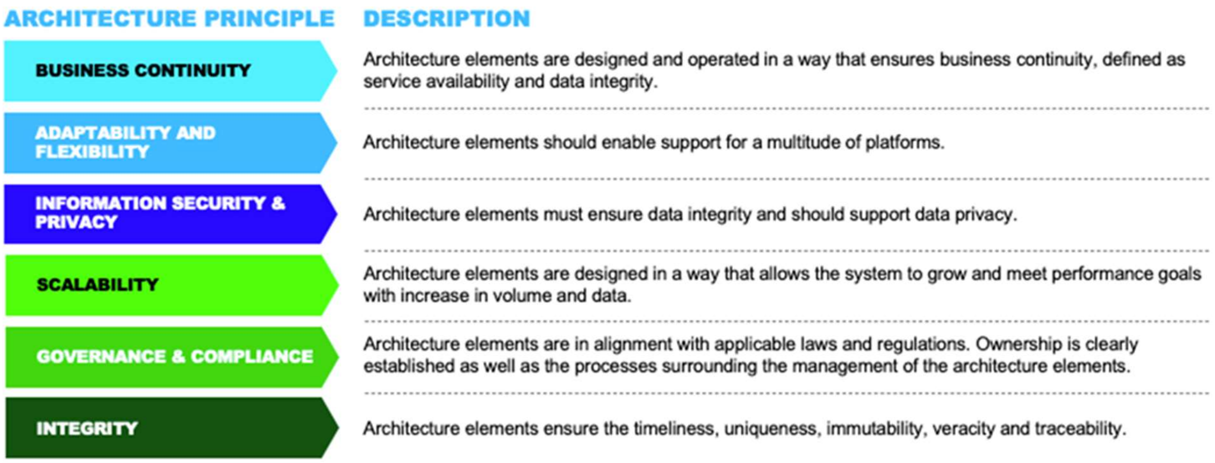

ARCHITECTURE PRINCIPLE	DESCRIPTION
BUSINESS CONTINUITY	Architecture elements are designed and operated in a way that ensures business continuity, defined as service availability and data integrity.
ADAPTABILITY AND FLEXIBILITY	Architecture elements should enable support for a multitude of platforms.
INFORMATION SECURITY & PRIVACY	Architecture elements must ensure data integrity and should support data privacy.
SCALABILITY	Architecture elements are designed in a way that allows the system to grow and meet performance goals with increase in volume and data.
GOVERNANCE & COMPLIANCE	Architecture elements are in alignment with applicable laws and regulations. Ownership is clearly established as well as the processes surrounding the management of the architecture elements.
INTEGRITY	Architecture elements ensure the timeliness, uniqueness, immutability, veracity and traceability.

Figure 22. Architecture Principles Framework

We based our design on the principles of adaptability and flexibility, information security and privacy, scalability, governance and compliance and integrity. As in the case of the non-functional requirement for performance, business continuity was assumed to be accomplished via the use of Kubernetes container-orchestration tools and blockchain's inherited architecture features such as fault-tolerance and immutability.

For adaptability and flexibility, we made the decision to use open standards for data exchange, specifically Verifiable Credentials 1.0 and Decentralized Identifiers 1.0. Additionally, infrastructure capabilities were selected to prevent vendor's locked-in. Separation of concerns was also a major decisional point, specifically in the way microservices were organized. Lastly, the use of human-centric design approaches was prioritized wherever possible.

Information security and privacy was a critical characteristic of the system design. Encryption protocols were established to secure cryptographic artifacts, as well as to protect the peer-to-peer communication channels. Allowing participants to completely own their cryptographic keys (including their generation) was also a priority. No personally identifiable information or any information considered confidential or highly sensitive was stored on the blockchain ledger by design. Additionally, in support of increased privacy, the use of Decentralized Identifiers minimizes the risk of correlation between transactions made in the blockchain network and the identity of their orchestrators. Finally, access rights were developed to be granularly managed, allowing admins to assign access privileges to control the creation, verification, modification, or deletion of data based on user role.

Although scalability was not a major requirement given the prototype scope, the system was designed to maximize production-readiness. Therefore, scalability principles were established to allow the architecture to easily scale across parties and the entire network. A cloud-first approach was used to facilitate and accelerate development. Additionally, several database systems were established for microservices to minimize the read and write operations against the blockchain ledger where not required, as well as to minimize the dependence on on-chain data to execute business logic. Moreover, the selection of Hyperledger Indy as the underlying blockchain platform provided a scalable consensus protocol satisfying conditions for byzantine fault tolerance (BFT) and performance.

For the governance and compliance, procedures were designed to on-board and off-board partners into the ecosystem. Additionally, several of the design decisions were based on pre-existing DoD guidelines and regulations for information security, including the required use of FedRAMP certified cloud service providers.

Finally, integrity was paramount, especially in a multi-party ecosystem of potentially disparate and external entities. Therefore, several key decisions surrounding data integrity were made to ensure the accuracy of the data processed and the information stored. Databases schemas were carefully designed to prevent data duplication and achieve a high level of information reciprocity between microservices. The use of blockchain technology already adds a higher degree of data integrity but does not protect against incorrect data provided by users ("Garbage In, Garbage Out"), therefore the system was designed to minimize the risk of storing incorrect information using data validation techniques on the presentation layer.

7.2 Solution Architecture

Figure 23 on the next page shows a high-level solution architecture diagram with details of the different components that were implemented within the prototype. Amazon Web Services (AWS) cloud was used to deploy all of the application components that were developed and the team leveraged some of AWS' out-of-the-box offerings to add enterprise-grade features in the areas of security, integration, data storage and infrastructure. In the following sections, each of the capabilities implemented within each service area will be described in more detail.

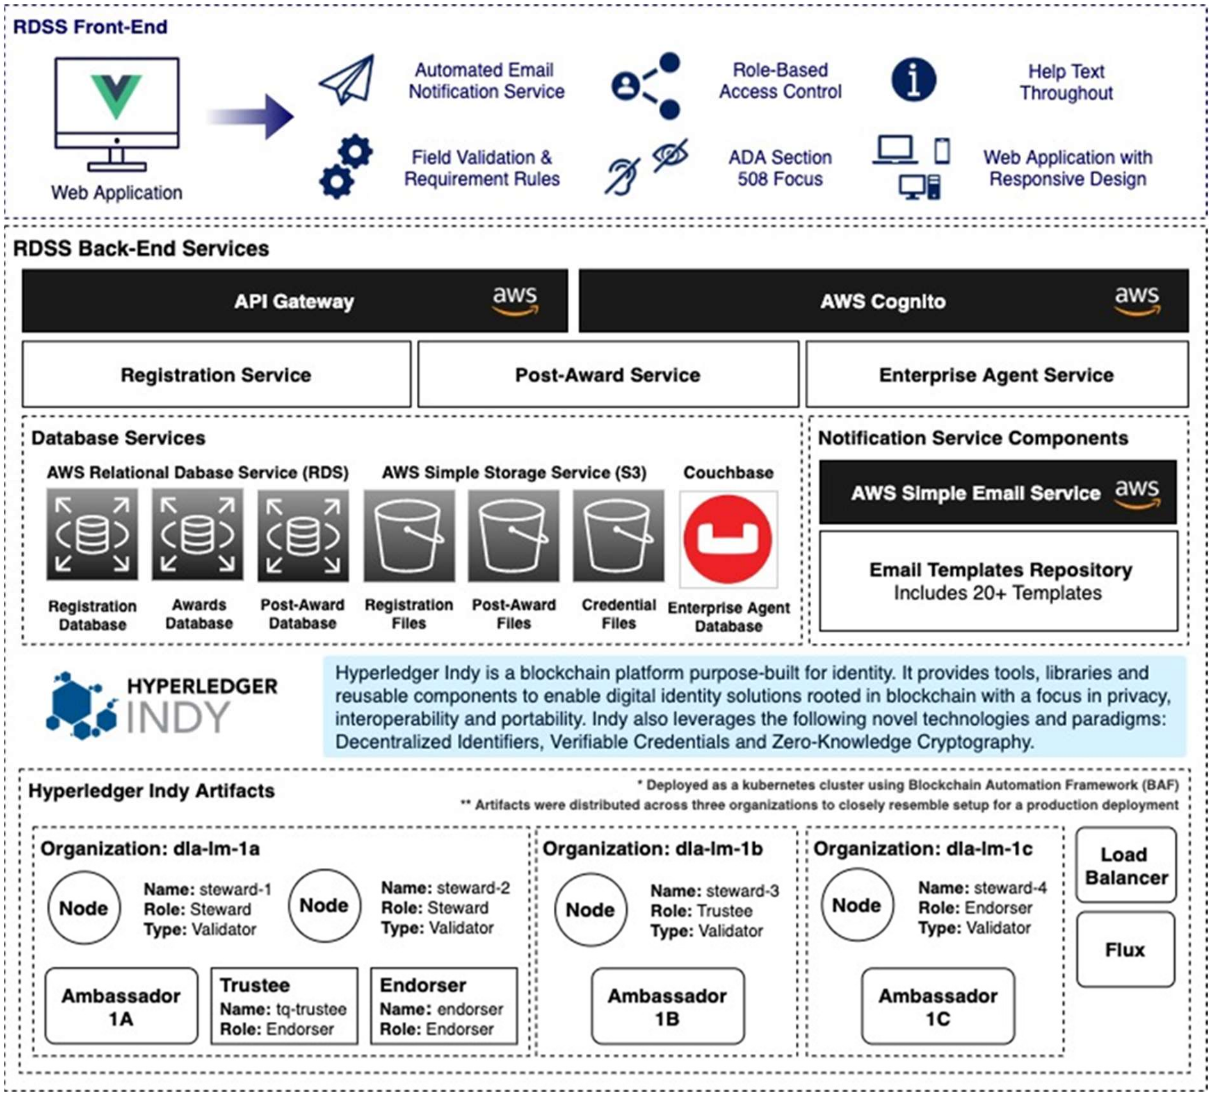

Hyperledger Indy is a blockchain platform purpose-built for identity. It provides tools, libraries and reusable components to enable digital identity solutions rooted in blockchain with a focus in privacy, interoperability and portability. Indy also leverages the following novel technologies and paradigms: Decentralized Identifiers, Verifiable Credentials and Zero-Knowledge Cryptography.

* Deployed as a kubernetes cluster using Blockchain Automation Framework (BAF)

** Artifacts were distributed across three organizations to closely resemble setup for a production deployment

Figure 23. High-Level Solution Architecture for Prototype Application

7.3 Security Services

Amazon Cognito was leveraged for user sign-up, sign-in and access control. Amazon Cognito is a fully managed service that provides a secure production-ready user directory and that allows setting up user pools without requiring additional server infrastructures. Additionally, Amazon Cognito is an enterprise-grade security module that offers multi-factor authentication and encryption of data-at-rest and in-transit.

Amazon Cognito enabled the implementation of a granular, role-based access control framework in the web application. Seven user pools were created to assign to registered users in the ecosystem. The roles and responsibilities for each user pool group are described as follows:

1. **Applicants -** A general user pool where newly registered users are assigned. When a CDAP Admin, in the case of newly registered companies, approves the company registration it assigned the registering user to the Vendor Admin user pool. If the company is already registered, then the corresponding Vendor Admin is the one

responsible for approving newly registered users, as well as assigning them either to the Vendor Admin or Vendor Standard user pools.

2. **CDAP-Admin -**User pool implemented for member of the CDAP team with admin responsibilities.
3. **CDAP-Standard -** User pool implemented for regular operators from the CDAP team.
4. **CDAP-Auditor -** User pool implemented for DLA roles that require read-only access privileges.
5. **Vendor-Admin -** User pool implemented for vendor employees with admin responsibilities.
6. **Vendor-Standard -** User pool implemented for vendor employees in charge to run general operations.
7. **Vendor-Auditor -** User pool implemented for vendor roles that require read-only access privileges.

As previously described, the web application supports mechanisms for external stakeholders to manage their company information, including their registered users access privileges. They are able to grant and/or revoke privileges autonomously, meaning, without requiring prior authorizations from CDAP or DLA in general. It is important to emphasize that external stakeholders are only granted the ability to manage access to application-level features. **Note:** They are not allowed to alter or directly manage access to access to AWS resources.

7.4 Data Services

This prototype solution leverages multiple data storage services provided by Amazon Web Services, including the Relational Database Service (or Amazon RDS) and Simple Storage Service (or Amazon S3). Amazon RDS offers scalable relational database capacity in addition to automating the execution of administrative tasks such as infrastructure setup, database configuration, patching and backups. All of these features are offered at very low rates for a cost-effective operation. AWS RDS was used to deploy the database services for the Registration and Award & Post-Award services, which are the core services supporting the Onboarding & Registration, and the Post-Award workflows respectively.

Additionally, a Couchbase cluster with two replicas was deployed as a containerized application within a Kubernetes cluster. The Couchbase database cluster was deployed as a convenient No-SQL off-ledger (or off-chain) data capability of the Enterprise Agent service.

7.5 Integration Services

This prototype solution implements three main integration services to handle the different business logic required to run the targeted workflows. Each service runs as an independent backend API, accessible via an instance of AWS API Gateway for authentication and access control. Additionally, these integration services were deployed as containerized applications hosted and managed within a Kubernetes cluster for security and resiliency. For more information regarding these services, please reference the technical Appendices provided separately.

The Registration Service is responsible of managing the registration of a new CAGE code in the CDAP program. Our design assumes that a company will be represented by a unique CAGE code, thus all demographic information, as well as users are tagged under the CAGE code of their corresponding employers. This service is the core component used throughout the Onboarding and Registration workflow. This workflow allows an authorized employee for a company in the ecosystem to register its company in the CDAP program. The registration follows a series of steps to create the company record on the database, enables a digital wallet to store credentials issued via the web application, and provisions admin access rights to the registering user for future account management activities or to onboard additional employees.

The Enterprise Agent Service provides capabilities to define, issue, store and verify verifiable digital credentials. It is composed by four sub-services: Controller, Communication, Aggregates and Agent.

1. **Controller:** This service handles the business logic that is specific to the type of credentials issued and used as part of the CDAP program. The Controller interacts with all the other components (Communication, Aggregates, and Agent) to execute the Credential Issuance and Credential Verification workflows.
2. **Communication:** This service is responsible for coordinating communication protocols with external agents.
3. **Aggregates:** This component handles the creation and management of data records stored in a Couchbase No-SQL database. We call this specific data service for the Enterprise Agent "off-chain" as it stores information of events that have occurred on the ecosystem outside of the distributed ledger. Only the owner of the Enterprise Agent instance has access to those records.
4. **Agent**: The core component of an Enterprise Agent implemented to handle all requests related to wallet operations (i.e. creating DIDs, storing credential, presenting credentials, etc.) and any other interactions with the nodes on the Hyperledger Indy network that are maintaining the distributed ledger.

The Award & Post-Award Service is where the business logic related to the CDAP program is handled. This includes simulating the award of a purchase order to a vendor, handling the submission of test reports and traceability documentation, including record keeping, verification, and processing. Moreover, the Award & Post-Award service provides additional features such as the automatic generation of Form 918 and the verification of the qualification credential claimed by the vendor. The service is also configured to send email notification to corresponding parties at different steps in the Post-Award workflow, with a final notification that serves as an official "Authorization to Ship" notification made to the vendor.

7.6 Blockchain Service

The underlying distributed ledger layer is running an instance of a Hyperledger Indy network with four validator nodes. The validator nodes are running images of an Indy-Node instance, which is a software application developed by the open-source community that supports Hyperledger Indy.

The validator node main responsibility is to maintain the distributed ledger. Some of the activities that nodes are responsible for are:

- Verifying, processing, and ordering transactions
- Verifying the identity and access roles of the agent submitting transactions
- Syncing the ledger with other nodes in the pool

- Participating in consensus protocols

The ledger consists of two parts: a log of all transactions and a Merkle tree. The transaction log is modeled as a sequence of key-value pairs, where a key is the sequence number of the transaction and value is the serialized transaction. The Merkle tree is a data structure built out of the hash of leaves and nodes. In the context of Hyperledger Indy, transactions act as leaves from which hashes are created to get child nodes up to the root node (or Merkle root). For this reason, a Hyperledger Indy ledger is considered a blockchain where each block contains only one transaction.

Each validator node replicates the ledger amongst all the other nodes in a pool and uses the Plenum consensus algorithm to keep all copies in-sync. Nodes are onboarded into a pool either by an entity that controls a node already authorized on the pool or via a genesis transaction that registers the node before deploying the network. Hyperledger Indy is a permissioned blockchain, thus the requirement to onboard nodes in a controlled way.

7.7 Infrastructure Service

Detailed information including the physical architecture deployed on AWS can be found within the separate technical Appendices. Other information about the infrastructure layer is listed below, which consisted of three virtual private clouds (VPCs):

1. **VPC – APIs:** Hosts a Kubernetes cluster using AWS Elastic Kubernetes Service (EKS), a classic load-balancer (CLB) ingress serving as a gatekeeper of the communication between the internet and the computing resources hosted within the cluster, and the master and worker nodes running the containerized applications for the Registration, Award & Post-Award, and the Agent Controller services.
2. **VPC – RDS:** Hosts the AWS Relational Database Services (RDS) for the corresponding APIs. Hosting these services in their own VPC increases security as they are not directly reachable from the public internet.
3. **VPC – Blockchain:** Hosts a Kubernetes cluster running the pool of Hyperledger Indy validator nodes.

Additionally, the following Amazon Web Services resources were utilized:

1. **AWS API Gateway:** An AWS resource used to manage hypertext transfer protocol (HTTP) calls routing from the web application to the corresponding services in the backend. Additionally, it secures all endpoints using JSON web token (JWT) authentication.
2. **AWS Route 53:** Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service. This service supports three main functions: domain registration, DNS routing, and health checking.
3. **AWS Lambda:** This AWS service offers an interface to execute code without implementing a complete backend server. This service was deployed to handle certain operations with AWS Cognito more easily.
4. **AWS Simple Email Service (SES):** This AWS service was implemented to enable the capability of sending email notifications. Preset email templates were developed that each backend API could dynamically populate and send to corresponding participants enabling need-to-know notifications and submission confirmations among other features.

7.8 Solution Strawman

Based on the information presented in previous sections, the various services and associated software packages and/or tools are shown in the solution strawman below. This software listed within the diagram should be evaluated against DLA Information Technology approved software packages to determine whether they are approved for use within pilot and/or production implementations.

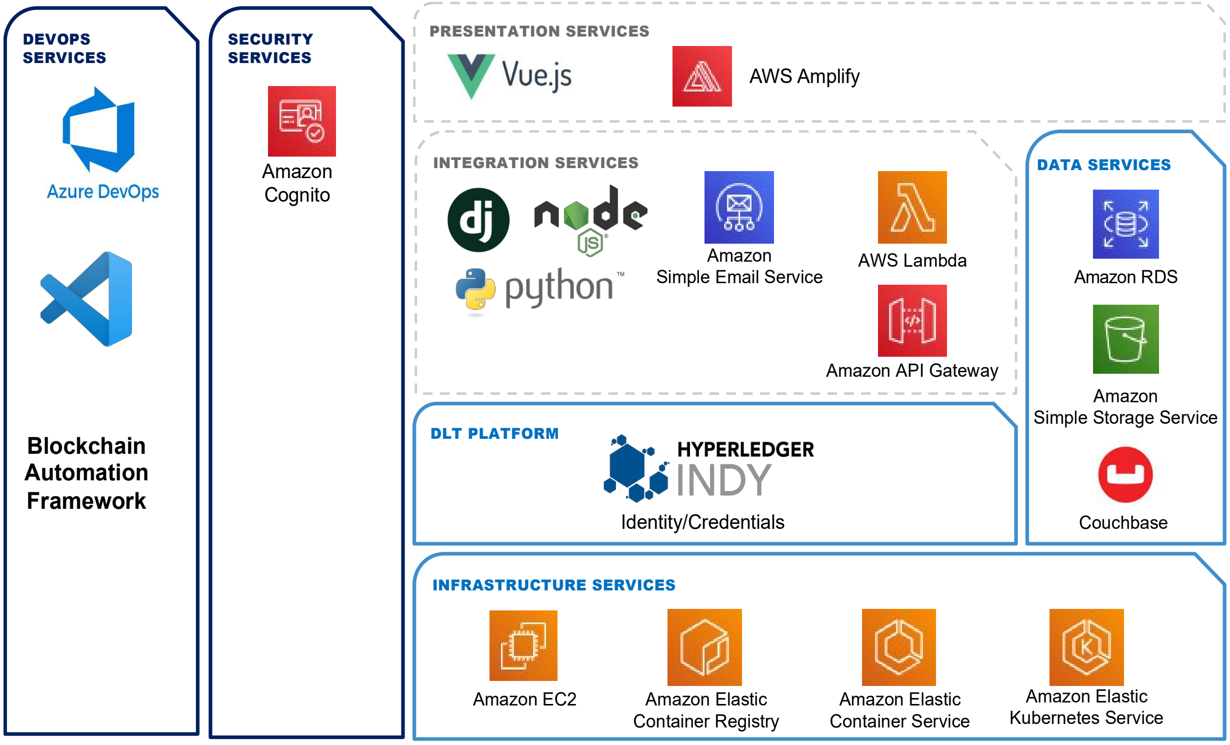

Figure 24. Prototype Solution Strawman Detailing Software and Resources Used

8.0 Prototype Application Features

Over the course of several weeks, the prototype application was developed using an Agile methodology and presentation of in-progress results to the Trusted Working Group on a bi-weekly basis. This process began with the development of the graphical user interface (GUI) using frontend development tools, followed by core microservices development once key functional aspects were confirmed. This section will describe the functionality and features of major workstreams within the application, but additional comprehensive details and user guides will be provided separately and/or in Appendices.

8.1 Vendor Functionality

The prototype application provides vendors with core capabilities that will allow them to complete all steps of the CDAP process. In general, these capabilities include (1) registering a company account, (2) applying for a qualification credential, and (3) submitting post-award documentation tied to an award as shown in the figure below.

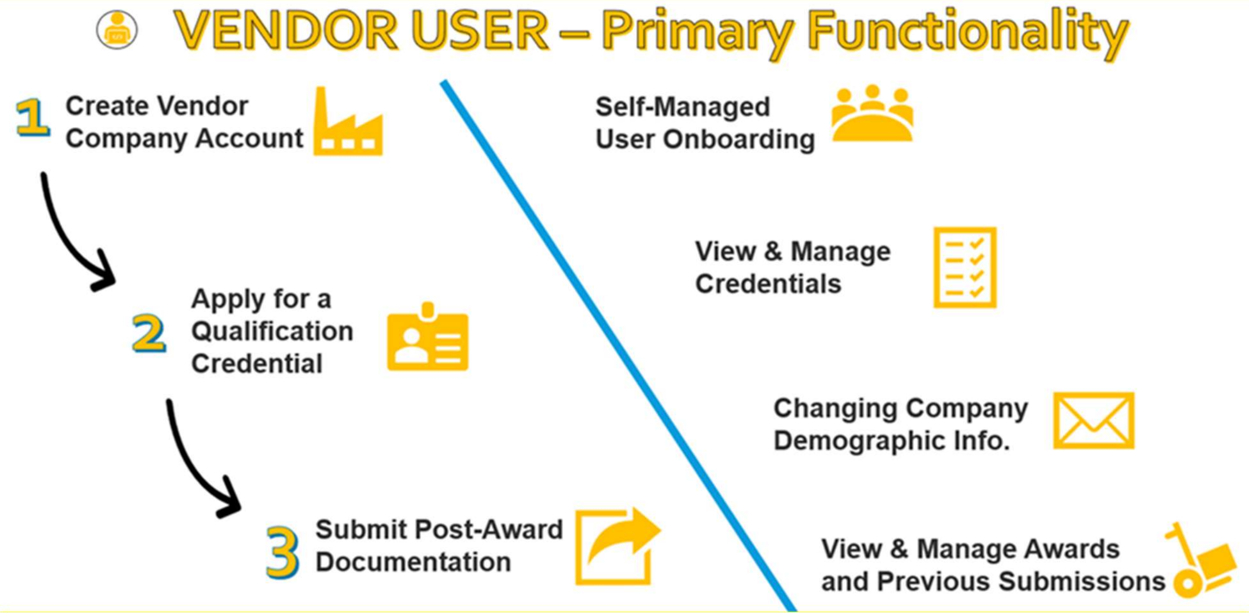
VENDOR USER – Primary Functionality
1 Create Vendor Company Account
2 Apply for a Qualification Credential
3 Submit Post-Award Documentation
Self-Managed User Onboarding
View & Manage Credentials
Changing Company Demographic Info.
View & Manage Awards and Previous Submissions

Figure 25. Primary Prototype Functionality and Features for Vendor Users

The functions listed above are made possible by a suite of pages provided within the Vendor (external) View of the prototype application. The major vendor functions are visible for all vendor-associated accounts, with some exceptions including the one-time DID registration and depending on the user role type. In general, the core functions for each vendor include:

1. **Apply for a Company Account** - The Company Account application includes both user and company information so that provisioning this account also provisions the company's first Administrative User account.
2. **Set Seed & DID to Onboard Blockchain** - Vendors set a secret Company Seed that will be used to create their Decentralized Identified (DID) on the CDAP blockchain. This completes the registration process.

3. **Apply for CDAP Credential** - Applying for a credential only requires the vendor to select the desired credential and attach proof/comment. Vendors are able to apply for and receive multiple credentials.
4. **Submit Post-Award Documentation** - Vendors will be notified of new awards. To submit Post-Award documentation, only four item information fields are required, along with any required traceability documentation.
5. **Ships Components to DLA** - Vendors will be able to track the status of all new and previous awards/submissions within the application. Vendors will also be notified via email when Authorization to Ship is given.
6. **Onboard New Users** - Vendors can onboard internal users, and the Role-Based Access Management allows for Administrators, Standard Users, and Auditors to share an application with differing views.
7. **Edit Company Information -** Vendors can view User & Company information on file, along with a view of all registered users for the company. Admins can submit legal changes of Company Address or Name.

8.2 CDAP Functionality

The prototype application provides CDAP users with functionality to manage all aspects of vendor onboarding and post-award traceability review. These capabilities mirror the vendor functions and generally include (1) onboarding new companies, (2) reviewing and issuing qualification credentials, and (3) reviewing post-award documentation to provide authorization to ship. These core functions are depicted in the figure below.

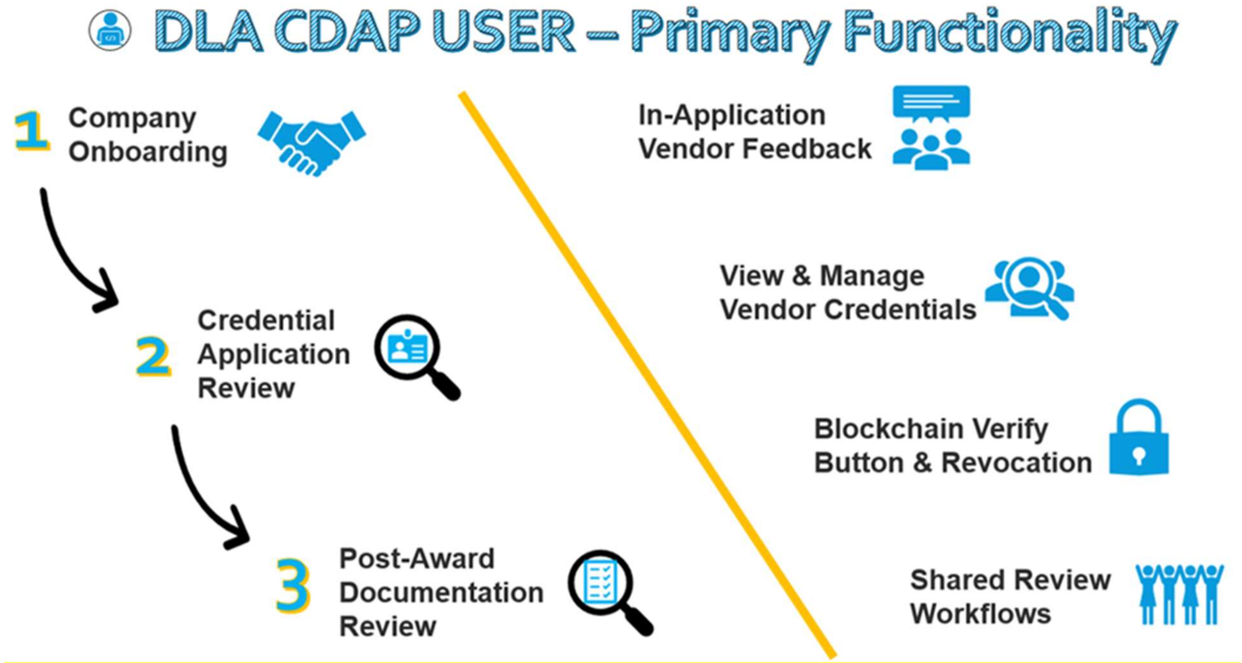
DLA CDAP USER – Primary Functionality
1 Company Onboarding
2 Credential Application Review
3 Post-Award Documentation Review
In-Application Vendor Feedback
View & Manage Vendor Credentials
Blockchain Verify Button & Revocation
Shared Review Workflows

Figure 26. Primary Prototype Functionality and Features for CDAP Users

The functions listed above generally mimic the vendor processes and provide DLA resources with a method for reviewing registration requests, credential requests, post-award submissions, and all related documentation. The system provides a process for returning comments to vendors, and automated emails alert any parties when a new message is received or action is triggered, so that actions can quickly be taken. In general, the core functions for the CDAP team include:

1. **Reviewing Company Account Applications -** Vendors must onboard by completing basic User & Company information forms. CDAP may review and be aware of incoming credential requests, and store data on file to re-use.
2. **Reviewing Credential Requests -** CDAP can review and store information submitted for Credential Requests. Applying is easy for vendors and they can be issued multiple credentials on-chain.
3. **Reviewing Post-Award Documentation -** Automated Form 918 generation, individual document review, vendor feedback, and dynamic statuses are a few of the features of the new Post-Award Documentation review.
4. **Viewing Vendor Data & User Data -** Core information about vendors is easily accessible to DLA users so that CDAP can drill-down into Credentials & Profile Information (including registered users) for any company.
5. **Onboarding New Users -** CDAP can onboard internal users, and the Role-Based Access Management allows for Administrators, Standard Users, and Auditors to share an application with differing views.
6. **Verifying Credentials -** Purpose-built using Hyperledger components, our blockchain solution allows for near real-time credential verification and revocation capabilities.

Much more detailed descriptions of each piece of functionality, all available prototype pages, and other relevant information can be found in the Application User Guides.

8.3 Primary Blockchain Features

The prototype application utilizes Hyperledger Indy to establish an identity-focused credential ledger for issuing, tracking, and revoking vendor credentials. The introduction of this structure in the application backend enables a series of unique capabilities that can enable greater collaboration and more efficient vendor documentation reviews. The primary features are listed below:

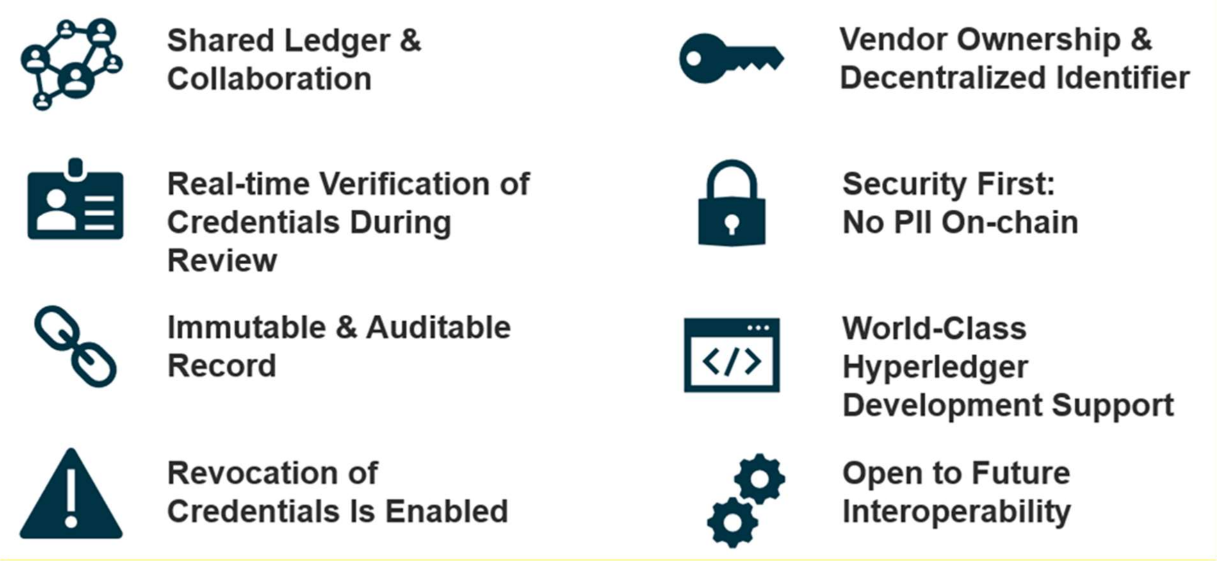
Shared Ledger & Collaboration
Vendor Ownership & Decentralized Identifier
Real-time Verification of Credentials During Review
Security First: No PII On-chain
Immutable & Auditable Record
World-Class Hyperledger Development Support
Revocation of Credentials Is Enabled
Open to Future Interoperability

Figure 27. Primary Prototype Blockchain Features

Some of the blockchain features, including development support tools, are a result of utilizing Hyperledger Indy, while others are best practices, inherent technical advantages, or features specifically built for this use case. A summary of the key features includes:

1. **Shared Ledger & Collaboration** – The shared blockchain ledger automatically enables better visibility for DLA and potential government partner stakeholders. This ledger provides credential information for all qualified vendors and can be viewed by inspecting company-level data.
2. **Real-Time Verification of Credentials During Review** – The near real-time 'Verify' Credential button, shown in the figure below, is a custom feature built for the CDAP team. This feature allows users to validate that vendor credentials have been issued and remain valid at any point throughout the review process. It is recommended that functional users check that vendor credentials are valid when making final 'Authorization to Ship' determinations. This button will also be very important if additional government stakeholders are brought on board to provide value and surety when verification of credentials issued by disparate organizations would otherwise be difficult.

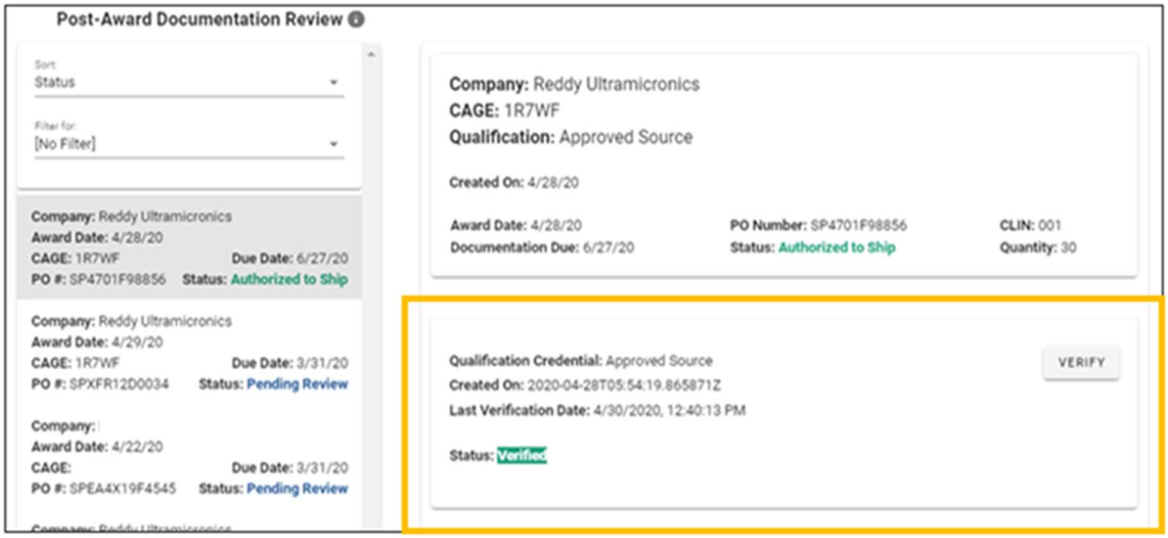

Post-Award Documentation Review

Sort: Status

Filter for: [No Filter]

Company: Reddy Ultramicronics
Award Date: 4/28/20
CAGE: 1R7WF Due Date: 6/27/20
PO #: SP4701F98856 Status: Authorized to Ship

Company: Reddy Ultramicronics
Award Date: 4/29/20
CAGE: 1R7WF Due Date: 3/31/20
PO #: SPXFR12D0034 Status: Pending Review

Company:
Award Date: 4/22/20
CAGE: Due Date: 3/31/20
PO #: SPEA4X19F4545 Status: Pending Review

Company: Reddy Ultramicronics

Company: Reddy Ultramicronics
CAGE: 1R7WF
Qualification: Approved Source

Created On: 4/28/20

Award Date: 4/28/20
Documentation Due: 6/27/20
PO Number: SP4701F98856
Status: Authorized to Ship
CLIN: 001
Quantity: 30

Qualification Credential: Approved Source
Created On: 2020-04-28T05:54:19.865871Z
Last Verification Date: 4/30/2020, 12:40:13 PM

Status: Verified

Figure 28. CDAP Post-Award Documentation Review Screen with Credential Verify Button

3. **No Personally Identifiable Information On-chain** – As a core security tenet, no personally identifiable information (PII) is stored on chain during the prototype or potential pilot phase. Although the blockchain transactions are secured via encryption and other cryptographic approaches, the blockchain transaction payloads should always avoid storing PII. Using the company identifiers to issue credentials avoids this issue and allows all PII to stay within DLA systems.
4. **Vendor Ownership & Decentralized Identifier** – When vendors are onboarded to the application, they are issued a vendor login and required to complete the process of registering a Decentralized Identifier (DID). This DID is a unique 32-character seed value that is used to create a unique ID for the company, similar to a CAGE code, to identify the company on the blockchain. Instructions for completing this process are sent to vendors

once their accounts have been provisioned and ensures that vendors are the 'owners' of their own data.

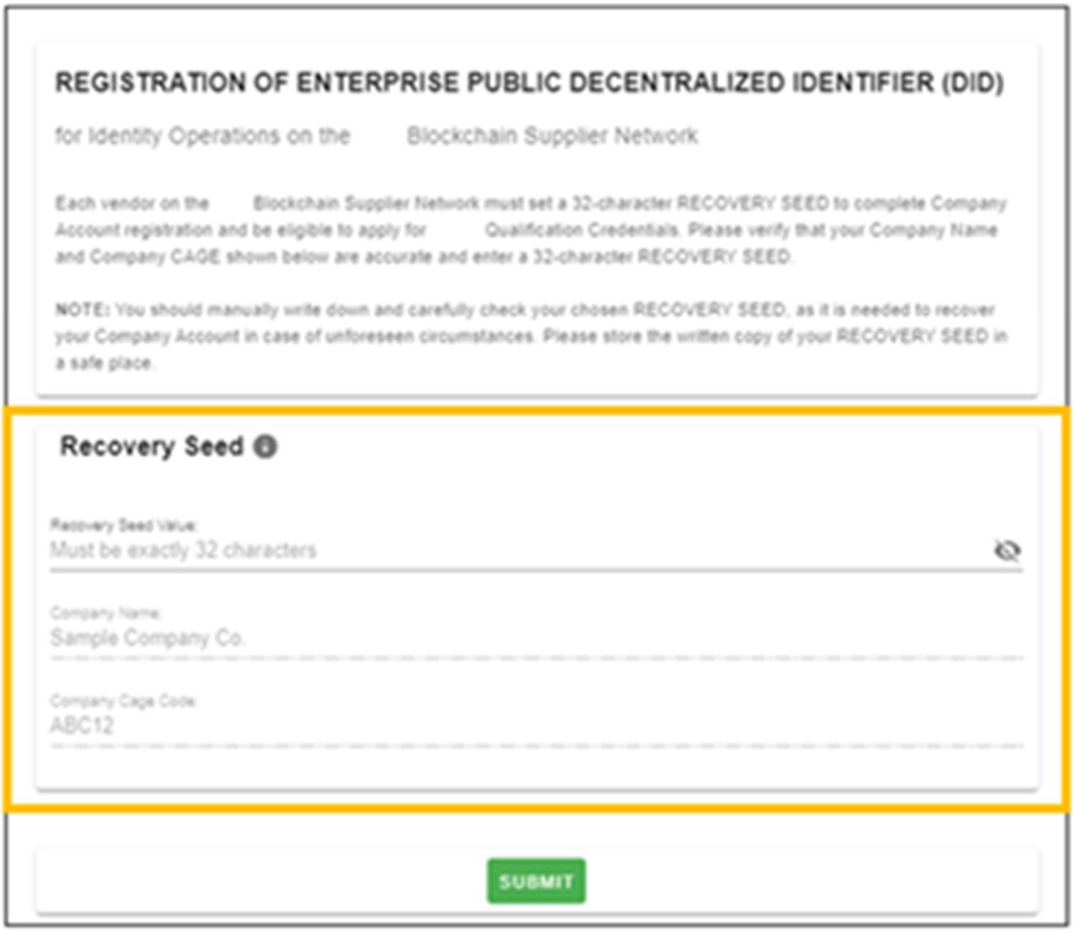

REGISTRATION OF ENTERPRISE PUBLIC DECENTRALIZED IDENTIFIER (DID)

for Identity Operations on the Blockchain Supplier Network

Each vendor on the Blockchain Supplier Network must set a 32-character RECOVERY SEED to complete Company Account registration and be eligible to apply for Qualification Credentials. Please verify that your Company Name and Company CAGE shown below are accurate and enter a 32-character RECOVERY SEED.

NOTE: You should manually write down and carefully check your chosen RECOVERY SEED, as it is needed to recover your Company Account in case of unforeseen circumstances. Please store the written copy of your RECOVERY SEED in a safe place.

Recovery Seed

Recovery Seed Value:
Must be exactly 32 characters

Company Name:
Sample Company Co.

Company Cage Code:
ABC12

Figure 29. Vendor Registration of Decentralized Identifier (DID) Page

5. **Immutable and Auditable Record** – Blockchain technology is immutable by nature and provides a high-level of auditability. Blockchain records are always added to existing blocks and never destroyed, so any 'changes' to a record such as a credential expiration are appended as new additions. This way, old records are preserved, but operational 'updated' information still reflects the current state. The figure below shows a record of a vendor's credentials and requests, including a rejected Supplier/Distributor credential and two active qualification credentials with attached details and documentation.

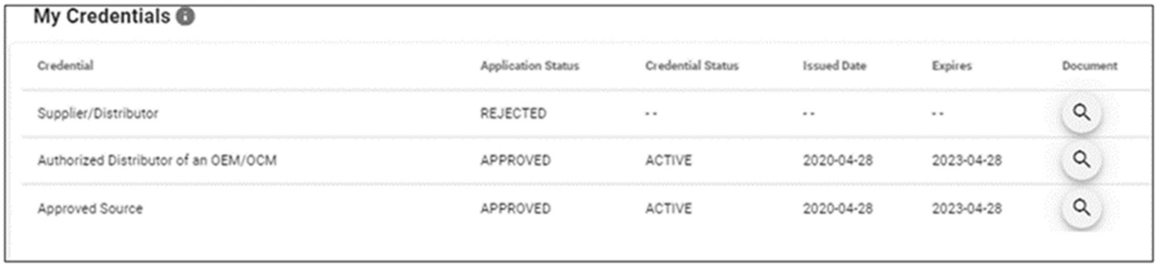

My Credentials

Credential	Application Status	Credential Status	Issued Date	Expires	Document
Supplier/Distributor	REJECTED	--	--	--	🔍
Authorized Distributor of an OEM/OCM	APPROVED	ACTIVE	2020-04-28	2023-04-28	🔍
Approved Source	APPROVED	ACTIVE	2020-04-28	2023-04-28	🔍

Figure 30. Vendor Credentials Page Showing Record of Credentials

6. **World-Class Hyperledger Development Support** – All Hyperledger products are backed by a large consortium of leading technical companies and other blockchain evangelists. With several thousand commits to the Hyperledger Indy code repository and a rich suite of development tools available, this platform is an excellent choice for identity-based solutions.
7. **Revocation of Credentials is Enabled** – One of the key features of Hyperledger Indy is the ability to revoke credentials. This is a key feature for any future-state where vendors and/or government partners may wish to issue their own qualification credentials such as 'Authorized Distributors (of an OEM)' or other government quality certifications. The ability to revoke a credential is crucial, because it allows each credential issuing group to maintain the right to revoke credentials at any time should new derogatory information come to light. This capability ensures that the credentials can be revoked in real-time and in-process procurements or post-award reviews will immediately reflect the change when checking with the 'Verify' Credential button.
8. **Open to Future Interoperability –** The blockchain application is separated into scalable microservices and built with potential future operability in mind. This means that the blockchain ledger is capable of sending and receiving information as needed to operate with other interfaces or business systems. The CDAP program collects item provenance information, so it might be a benefit to DLA to eventually stand up a parallel item track and track blockchain, while this identity blockchain ties those records to a vendor and active verifiable credential.

8.4 Primary Application Features

The prototype application provides many inherent advantages by utilizing blockchain technology. However, there are several other functional issues and pain points that our application attempted to solve for with other conventional development techniques. A summary of the primary application features are listed below.

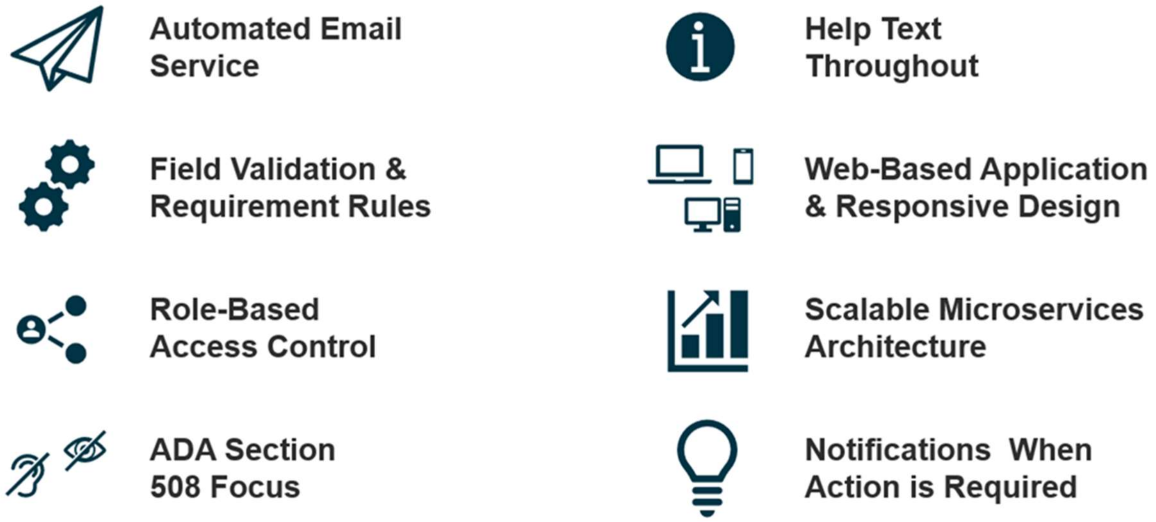
Automated Email Service
Help Text Throughout
Field Validation & Requirement Rules
Web-Based Application & Responsive Design
Role-Based Access Control
Scalable Microservices Architecture
ADA Section 508 Focus
Notifications When Action is Required

Figure 31. Primary Application Features for the Prototype

Many of these application features include time-saving functions that will ensure a more efficient CDAP review process and a higher quality of incoming data. Other features were included to make the application more user friendly and to ensure that both vendors and DLA users find the application intuitive. Each of these features are described in more detail below.

- **Automated Email Service** – An automated email service performs several functions within this prototype application. The email service provides an 'email authentication' function that prevents potential applicants from spamming DLA with registration requests, by requiring verification of email before the requests process. In addition, email messages are sent to confirm all major form confirmations and to indicate when any new actions are required to keep reviews flowing. A detailed listing of all automated email scenarios is included within the ***Application User Guide***.
- **Field Validation & Requirement Rules** – Anywhere that users are required to enter data, field validation rules were introduced (if a standard format is required) to ensure new data matches expected formats. This will help ensure data quality by preventing obviously erroneous entries and allowing users to identify mistakes before they submit data for review. This feature should help improve data quality, reduce error rates, and reduce rejections overall.
- **Role-Based Access Control** – The application provides four different 'primary' user roles with varying levels of page visibility, access, and control. The Universal Admin role is only available to Core CDAP Team members and meant to only be used in extraordinary circumstances to solve technical issues. The Administrative roles are meant to provide both DLA and vendors with a high-level of functional control within their domains (i.e. with regard to DLA or vendor functions), and manage sensitive tasks such as changing company address information or user onboarding (including new user role assignment). The Standard User role is meant to perform most day-to-day functions within both organizations and can facilitate primary CDAP submissions and reviews. Finally, the Auditor role exists as a 'read-only' type of role with limited access and visibility meant for ancillary support resources within both organizations who do not actively participate in CDAP exchanges.
- A**DA Section 508 Focus** – Where possible, aspects of the Americans with Disabilities Act (ADA) Section 508 accessibility guidance were incorporated into the application design. Color palettes selected to indicate 'success/approval' or 'failure/rejection' were created with colorblindness in mind, a plethora of instructions are provided throughout, and the beginnings of code-based accessibility were developed to conform with applications such as e-readers or other tools in future phases.
- **Help Text Throughout** – The entire application features help text (accessed via 'info' buttons) throughout every step and page. The help text provides information about the purposes of each page, the definitions for each required field, and other contextual help such as qualification definitions when a vendor is applying for a new qualification credential. This help text should help answer questions that may arise during the use of the application and reduce the number of errors or incomplete data submitted to DLA.
- **Web-Based Application & Responsive Design** – The prototype application was developed on the web to provide access to all user groups, without the need for special

networking. In addition, the application was built to be responsive – the pages and views can shift and scale differently according to the size of the monitor or web browser window present. This will enable users to view the application well regardless of their viewing platform and can enable users to customize their preferred viewing dimensions by expanding or contracting their browser window.

- **Scalable Microservices Architecture** – The prototype application was built with individual microservices for handling major workflows such as registration, credentialing, and post-award submission. Each microservice was built independently so that it can scale and handle higher workloads and data interfaces in a potential pilot or production phase. Although this design was more time-intensive, it will save time and effort in future phases and ensure scalability.
- **Notifications When Action is Required** – Notifications are provided to end users (or shared workload emails) whenever an action is required by either DLA or vendors. This process will ensure that reviews are completed in a timely manner and should help greatly reduce the current review time.

9.0 Benefit Cost Analysis

The goal of the Blockchain Traceability for the Counterfeit Detection & Avoidance Program was to build a prototype application based on blockchain technology and confirm the business benefits of such a capability. However, now that the prototype application has been successfully developed, demonstrated, and tested by end users, a business case must be made for implementing this capability. As a supply chain risk management (SCRM) and anti-fraud focused initiative not inherently based on direct financial savings, it will be inherently difficult to measure a direct quantitative benefit for this solution. However, numerous qualitative benefits may come from such a capability and potential indirect cost savings may exist where efficiencies are introduced.

9.1 Benefit Cost Analysis Background

The basic approach for any benefit-cost analysis (sometimes also referred to as a Business Case Analysis – both referred to as 'BCA') includes understanding the costs associated with implementing a capability and weighing those against potential savings. These costs can come in multiple forms (one-time or recurring costs) and can be compared to multiple forms of benefits including qualitative and quantitative types. This analysis typically considers the 'Cost-Benefit ratio' for the current situation, or status quo, against the scenario(s) where a change is made, which in this case would be the implementation of new capability such as the Blockchain Traceability for CDAP application.

Costs may include any monetary or non-monetary impediments required or caused by the implementation of a solution. The two primary types are one-time costs and recurring costs, as defined below.

- **One-Time Costs:** Costs that are incurred only once. (e.g. Blockchain Application production build and integration)
- **Recurring Costs:** Costs that repeat on some interval. (e.g. hosting, cloud, or other ongoing costs)

Benefits may include any monetary or non-monetary gains which are realized as a result of the solution implementation. The two primary types are tangible benefits, which come in many forms, and intangible benefits.

- **Tangible Benefits:** Quantifiable benefits which can be measured or assessed, including both cost savings and cost avoidance.
    - **Cost Savings:** Any action that results in a tangible financial benefit that lowers current spending, investment, or debt levels. (e.g. making the current processes more efficient so that CDAP reviewers are able to complete more reviews in the same amount of time).
    - **Cost Avoidance:** Any action that avoids having to incur costs in the future (e.g. introducing data controls and validations to reduce the number of backorders caused by errors).
- **Intangible Benefits:** Non-quantifiable but significant qualitative improvements or preventative measures. (e.g. the real-time verification button ensures that reviewers can validate qualification validity at the time of review and halt reviews when a credential has been revoked or is invalid.)

9.2 Qualitative Benefit

The Blockchain Traceability for CDAP application is meant to primarily improve the business functions of the Counterfeit Detection and Avoidance Program, and provide technical enhancements that will improve the veracity and legitimacy of vendor credentials exchanged during the repeatable post-award review process. As a tool focused on preventing fraud and facilitating a more efficient review of traceability, test reports, and quality-related items, there are a substantial amount of qualitative (non-measurable, intangible) benefits to be realized by implementing such a solution. A summary of some of the major qualitative benefits for this salutation are listed below and should be carefully considered along with quantitative benefits presented in the next section.

- The real-time verify feature will provide Core CDAP users (and potentially Procurement users) the ability to ensure vendors are still qualified at the appropriate level during post-award reviews. This provides a level of credential surety that was not previously possible with paper/PDF copies provided for verification.
- The systemic digital storage of qualification documentation at the time of application (documents proving qualification level) will ensure full documentation is available and auditable for future review to backup qualification decisions or protests at any time.
- The systemic digital storage of award-specific post-award documentation and traceability will ensure information is available to backup 'Approval/Rejection' decisions for any protests related to reviews.
- DLA can combine collected Traceability Documentation with existing datasets to provide increased knowledge of supplier relationships and enhance other SCRM capabilities with this data.
- Shared workloads will increase efficiency and awareness for both vendors and CDAP, allowing more timely reviews and lessening the likelihood of a 'lost review'
- Comprehensive help text provided throughout the application will reduce the vendor's reliance on the CDAP team to complete forms and submissions, and will provide vendors with information on a need-to-know basis as they work through registration, qualification, and post-award documentation submission processes.
- Providing notification emails and confirmations when forms are submitted will introduce a level of automation that will provide peace of mind to vendors and enable more efficient transactions. Less time will be wasted waiting for the next step of the process since stakeholders will be notified any time an action or review is required.
- The enhanced data validation rules on every data entry form and reduced number of required fields will drastically reduce the number of data errors made by vendors and the resultant backorders, re-procurements, and material availability issues that can result from those errors.
- The automation of key portions of this workflow will allow Core CDAP reviewers to focus on reviewing Post-Award documentation, test reports, and traceability instead of responding to administrative questions and manually entering vendor-submitted PDFs.
- Data will be preserved and immutable lending itself to other applications
- Increased data collection will allow for tracking of new metrics including supplier performance metrics and quality metrics.

- Increased workload awareness will provide efficiencies to various component departments by allowing CDAP components views of incoming reviews and an understanding of expected future workload.

9.3 Quantitative Benefit

Many DoD and external reports and investigations have served as a way to estimate the costs associated with not implementing supply chain risk management (SCRM) capabilities once major fraud activities have occurred. Unfortunately, this retroactive estimation is the only reliable approach available for estimating the magnitude of potential costs when a catastrophe and/or significant loss occurs. As supply chain risk management capabilities such as this blockchain prototype are not strictly cost-saving measures, their quantitative benefit can primarily be estimated by conservative estimates regarding avoided costs associated with a small number of compromised components. This section will attempt to provide conservative baseline estimates for a status quo situation within DLA and compare those to potential future state options for a blockchain solution based on expected awards processed (summarized by FSC, to gauge throughput) and general implementation complexity (provided at 3 complexity levels to estimate cost).

9.3.1 General Model for Quantitative Benefits

To perform the economic analysis of quantitative benefits, we used guidelines from DoD and the Office of Management and Budget (OMB) to evaluate two scenarios against the current-state CDAP review process (as a baseline): 1) Blockchain for FSC 5962 awards and 2) Blockchain for FSC 5961, 5962, 5963, 5998, and 5999 awards. Calculations for labor costs, including direct labor, labor associated to error handling and labor associated to re-procurement were done for the Baseline and the two Blockchain scenarios. These calculations were then compared to determine whether the Blockchain solution provides opportunities for cost-savings or cost-reduction, which will correspond to the quantitative benefits. Additional cost estimates were also considered for the Blockchain scenarios to account for required sustainment support, such as Tech Support once the solution is implemented, and recurrent cost associated to hosting the application on the Cloud.

Based on the assumptions provided on the next page, expected cost savings and/or avoidances were estimated in the following major areas:

- **Post-Award Cost Reduction** – Reduction in the amount of time required to analyze documentation and paperwork provided by suppliers. This includes a reduction in the amount of time required corresponding to vendors, manually entering information, and generally doing other tasks not directly tied to reviewing traceability and/or relevant quality information. The reduction in time will reduce the amount of labor cost incurred to review awards.
- **Re-procurement Cost Avoidance** - Repurchasing items after they are found to be defective or re-soliciting a procurement due to duplicitous errors is a problem within the CDAP process. When DLA finds a non-conforming or potentially malicious/counterfeit component, it is forced to create a new requisition and go through the procurement process again to acquire the item. This application will reduce the likelihood a malicious or non-conforming component is procured by ensuring only trusted vendors with verifiable credentials are included within the program.

Table 18. Assumptions for Quantitative Benefit Calculations

General	
Total Annual Hours per FTE	1,920
Discount Rate for Present Value Computation	7%
Total Periods for Financial Analysis (Years)	10
Average Cost per Contract	$14,500.00
Salary Rates ($ per Hour)	
Reviewer	$37.42
CDAP Tech Quality / Test Center	$57.49
Product Specialist	$41.61
Contracting Officer	$49.48
Developer	$37.42
DevOps Specialist	$50.00
Awards by FSC	
Total Awards in a Year	11,480
FSC 5962	1,483
FSC 5961	1,744
FSC 5963	289
FSC 5998	3,101
FSC 5999	4,863
Baseline - CDAP Operation	
Error Rate	39%
Re-procurement Rate	5%
Re-procurement Operation	
Total Cost per Re-procurement	$838
1 Product Specialist FTE ~ 15 Hours per Re-procurement	$624
1 Contracting Officer FTE ~ 2 Hours per Re-procurement	$99
1 CDAP FTE ~ 2 Hours per Re-procurement	$115
Future State - Blockchain Solution	
Error Rate	1%
Re-procurement Rate	0%
Implementation Cost	
Low Complexity	$2,500,000
Medium Complexity	$3,500,000
High Complexity	$4,500,000
Sustainment Costs	
Tech Support Team	
Total Cost Tech Support per Year	$95,846
2 Developers FTEs ~ 960 Hours per Year	$71,846
1 DevOps Specialist FTE ~ 480 Hours per Year	$24,000
Cloud Hosting Costs ($ per month)	$2,000

9.3.1.1 Definition of Calculations

- **Total Potential Re-procurements in a Year -** This represents the number of potential awards that might require re-procurement.
- **Direct Labor Cost -** These are the costs to review an award when no errors were detected.
- **Labor Associated to Error Handling** - These are the labor costs related to re-evaluating the awards where errors were found. Some of these tasks include contacting the vendors so that they resubmit documentation, and any other administrative tasks such as re-uploading information into the CDAP Catalog.
- **Labor Associated to Re-procurement** – This refers to the labor cost incurred while re-procuring parts. As listed in the List of Assumptions, it estimated that each re-procurement requires at least 15 hours for a Product Specialist to evaluate the part and find potential substitutions, 2 hours for a Contracting Officer to manage the award of the purchase order and 2 hours associated to members of the CDAP team that are tasked to provide more in-depth reviews of the part's quality requirements. These estimated labor allocations are translated into costs using assumed salary rates for each actor in the process.

9.3.1.2 Baseline Scenario

The Baseline scenario represents the current CDAP review process. For this scenario, calculations are performed for direct labor, additional labor required to handle errors and associated labor to re-procure parts when awards get cancelled due to failed reviews[1].

Table 19. Time Estimates for CDAP Tasks in Various States

CDAP - Current State Tasks	Clean State	Error State
Review Traceability	30	60
Communication with Vendor	3	10
Product Specialist	2	3
Upload Data to CDAP Catalog	10	10
Sign & Send Shipment Notification	7	7
Total Time (minutes)	**52**	**90**

Estimated time measurements, shown in the table above are used to perform each labor cost calculations for general reviews, error handling and re-procurement. For error handling, only the difference in time between Clean and Error state is used to avoid double counting. That can be interpreted as the additional time it takes to handle errors.

Calculations for Only FSC 5962:

- **Total Potential Re-procurements in a Year** = (# of Awards in a Year **x** Re-procurement Rate) = (1,483 awards per year **x** 5%) = **44 awards per year**

[1] A failed review is a review that did not satisfy post-award documentation requirements appropriately after several resubmissions.

- **Direct Labor Cost** = (Awards in Year **x** Clean State Time **x** Salary Rate Reviewer) = (1,483 awards per year x 0.86667 hour per award x $37.42 per hour) = **$48,095**
- **Labor Associated with Error Handling** = (Awards in Year **x** Additional Time for Errors **x** Salary Rate Reviewer) = (1,483 awards per year **x** 0.63333 hours per award **x** $37.42 per hour) = **$35,146**
- **Labor Associated with Re-procurement** = (Cost per Re-procurement **x** Total Potential Re-procurements in a Year) = ($840 per award **x** 44 awards per year) = **$37,372**

Calculations for All 5 FSCs (5961, 5962, 5963, 5998, and 5999):

- **Total Potential Re-procurements in a Year** = (# of Awards in a Year **x** Re-procurement Rate) **=** (11,480 awards per year x 5%) = **344 awards per year**
- **Direct Labor Cost** = (Awards in Year x Clean State Time x Salary Rate Reviewer) = (11,480 awards per year **x** 0.86667 hour per award **x** $37.42 per hour) = **$470,177**
- **Labor Associated to Error Handling** = (Awards in Year **x** Additional Time for Errors **x** Salary Rate Reviewer) = (11,480 awards per year **x** 0.63333 hours per award **x** $37.42 per hour) = **$272,068**
- **Labor Associated with Re-procurement** = (Cost per Re-procurement **x** Total Potential Re-procurements in a Year) = ($840 per award **x** 344 awards per year) = **$289,296**

Note: All dollar amounts have been rounded to the next integer to simplify by avoiding cents.

9.3.1.3 Blockchain Scenarios

For Blockchain Scenarios, calculations are performed for direct labor and additional labor required to handle errors. Labor associated to re-procure parts when awards get cancelled due to failed reviews[2] are assumed to be null (or zero) because the implemented controls in the solution are expected to help eliminate the need to re-procure parts due to errors on the submitted documentation. This has a great impact in terms of associated costs when compared to the Baseline scenario, as costs associated to re-procure parts are avoided.

Time Estimates for CDAP Tasks in Various States

CDAP - Current State Tasks	**Clean State**	**Error State**
Review Traceability	30	30
Communication with Vendor	0	5
Product Specialist	0	0
Upload Data to CDAP Catalog	0	0
Sign & Send Shipment Notification	0	0
Total Time (minutes)	**30**	**35**

The blockchain solution automates the communication with vendors, reducing the time for the tasks *Communication with Vendor* to approximately 5 minutes. *Sign & Send Shipment Notificatio*n task is assumed to be zero because the Blockchain solution implements automatic emails for

[2] A failed review is a review that did not satisfied post-award documentation requirements appropriately after several resubmissions.

Shipment Notification. Additionally, *Upload Data to CDAP Catalog* is also assumed to be zero as this is done as soon as the post-award documentation and Form 918 are submitted via the web application.

Implementing the blockchain solution will have a cost in the form of a one-time initial investment. Furthermore, the solution will require sustainment support, which represents additional recurrent costs to the operation. Sustainment support costs were assumed to include a dedicated Tech Support Team and cost related to hosting the application on the Cloud. Details for Sustainment Support are shown in the table below.

Table 20. Sustainment Support Costs

Sustainment Support Resource	FTEs	Annual Hours	Total Annual Cost
Developer	2	960	$71,846
DevOps Specialist	1	480	$24,000
Tech Support			**$95,846**
Cloud Hosting			**$24,000**
Total Sustainment Cost			**$119,846**

The most significant cost associated with implementing a blockchain solution is the one-time implementation cost. Implementation costs were estimated by the approximate level of complexity to deploy the solution and are shown in the figure below. **NOTE:** these figures are simply provided for rough approximation purposes and should not be taken to constitute an official quote for work.

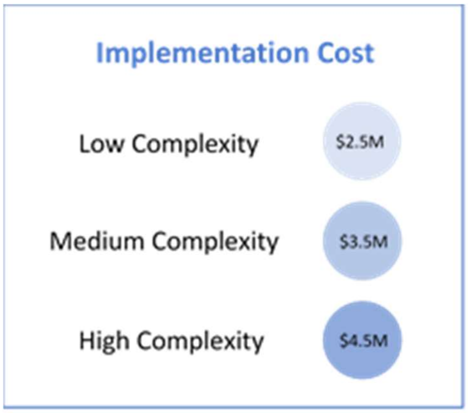

Figure 32. Rough Estimates of One-Time Implementation Cost by Complexity Level

Calculations for Only FSC 5962:

- **Direct Labor Cost** = (Awards in Year **x** Clean State Time **x** Salary Rate Reviewer) = (1,483 awards per year **x** 0.5 hours per award **x** $37.42 per hour) = **$27,747**
- **Labor Associated with Error Handling** = (Awards in Year **x** Additional Time for Errors **x** Salary Rate Reviewer) = (1,483 awards per year **x** 0.08333 hours per award **x** $37.42 per hour) = **$4,624**

Note: All dollar amounts have been rounded to the next integer to simplify by avoiding cents.

Calculations for All 5 FSCs:

- **Direct Labor Cost** = (Awards in Year **x** Clean State Time **x** Salary Rate Reviewer) = (11,480 awards per year **x** 0.5 hour per award **x** $37.42 per hour) = **$214,791**
- **Labor Associated with Error Handling** = (Awards in Year **x** Additional Time for Errors **x** Salary Rate Reviewer) = (11,480 awards per year **x** 0.08333 hours per award **x** $37.42 per hour) = **$35,798**

Note: All dollar amounts have been rounded to the next integer to simplify by avoiding cents.

9.3.1.4 Evaluation of Alternatives - FSC 5962 Only

This scenario calculates the proposed benefit of utilizing blockchain to facilitate the current workload of FSC 5962 awards only through the Counterfeit Detection and Avoidance Program. As these calculations will show, it may not be financially feasible to implement a blockchain-based solution for CDAP if the workload remains constant with only FSC 5962 awards. The alternate option of increasing the FSC workload to include other at-risk component classes (5961, 5962, 5963, 5998, and 5998) will also be evaluated.

Calculations of Benefits:

- **Cost-Avoided of Labor Associated with Re-procurement**

Labor/Re-procurement Cost-Avoided = (Baseline – Blockchain) = ($37,372 - $0) = **$37,372**
Present Value = $37,372 **x** (P/A,7%,10 years) = **$262,482**

- **Cost-Reduction of Labor Associated with Error Handling**

Labor/Error Handling Cost-Reduction = (Baseline – Blockchain) = ($35,146 - $4,624) = **$30,522**
Present Value = $30,522 **x** (P/A,7%,10 years) = **$214,371**

- **Cost-Reduction of Direct Labor**

Labor/Direct Labor Cost-Reduction = (Baseline – Blockchain) = ($48,095 - $27,747) = **$20,348**
Present Value = $20,348 **x** (P/A,7%,10 years) = **$142,914**

- **Present Value of Total Benefits**

PV of Total Benefits = $262,482 + $214,371 + $142,914 = **$619,768**

Calculations of Costs Associated with Blockchain Solution:
Implementation Costs are already assumed to be done at present time; thus, no calculations are required other than present value adjustment.

- **Present Value of Tech Support Labor Costs**

  PV of Tech Support Team Costs = $95,846 **x** (P/A, 7%, 10 years) = **$673,185**

- **Present Value of Cloud Hosting Cost**

  PV of Cloud Hosting Cost = $24,000 **x** (P/A, 7%, 10 years) = **$168,566**

- **Present Value of Total Costs for Low Complexity**

  PV of Total Costs = $673,185 + $168,566 + $2,500,000 = **$3,341,751**

- **Present Value of Total Costs for Medium Complexity**

  PV of Total Costs = $673,185 + $168,566 + $3,500,000 = **$4,341,751**

- **Present Value of Total Costs for High Complexity**

  PV of Total Costs = $673,185 + $168,566 + $4,500,000 = **$5,341,751**

Low Complexity Analysis:

- **NPV** = Total Benefits at Present Value – Total Costs at Present Value = $619,768 - $3,341,751 = **($2,720,359)**
- **Benefit-Cost Ratio** = Total Benefits at Present Value / Total Costs at Present Value = $619,768 / $3,341,751 = **0.19**
- **Payback Period** = Implementation Cost / Net Annual Cash Flow = Implementation Cost / abs(Annual Benefits – Annual Costs) = $2,500,000 / abs(($37,372 + $30,522 + $20,348) – ($95,846 + $24,000)) = $2,500,000 / $31,374 = **80 years or > 20 years**

Medium Complexity Analysis:

- NPV = Total Benefits at Present Value – Total Costs at Present Value = $619,768 - $4,341,751 = **($3,720,359)**
- Benefit-Cost Ratio = Total Benefits at Present Value / Total Costs at Present Value = $619,768 / $4,341,751 = **0.14**
- **Payback Period** = Implementation Cost / Net Annual Cash Flow = Implementation Cost / abs(Annual Benefits – Annual Costs) = $3,500,000 / abs(($37,372 + $30,522 + $20,348) – ($95,846 + $24,000)) = $3,500,000 / $31,374 = **111 years or > 20 years**

High Complexity Analysis:

- **NPV** = Total Benefits at Present Value – Total Costs at Present Value = $619,768 - $5,341,751 = **($4,720,359)**
- **Benefit-Cost Ratio** = Total Benefits at Present Value / Total Costs at Present Value = $619,768 / $5,341,751 = **0.12**

- **Payback Period** = Implementation Cost / Net Annual Cash Flow = Implementation Cost / abs(Annual Benefits – Annual Costs) = $4,500,000 / abs(($37,372 + $30,522 + $20,348) – ($95,846 + $24,000)) = $4,500,000 / $31,374 = **142 years or > 20 years**

Summary of Evaluation – FSC 5962 Only

Based on an analysis of the potential costs associated with a blockchain implementation, it may not be feasible to do so if it is used solely for awards within FSC 5962. This is due in large part to the fact that this supply class only includes about 1,500 procurements per year and would not generally provide enough benefit to cover the expected costs within a reasonable timeframe. The NPV is negative for all complexity options and BCRs below 1 indicate financial infeasibility.

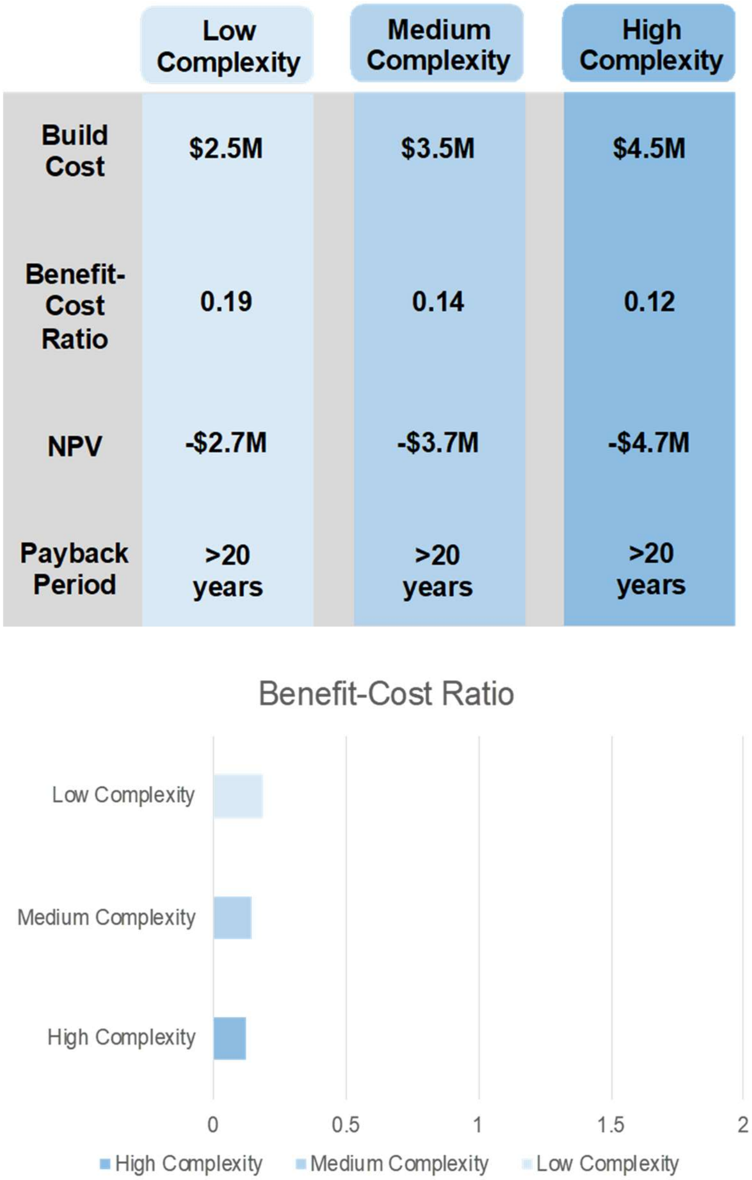

	Low Complexity	Medium Complexity	High Complexity
Build Cost	$2.5M	$3.5M	$4.5M
Benefit-Cost Ratio	0.19	0.14	0.12
NPV	-$2.7M	-$3.7M	-$4.7M
Payback Period	>20 years	>20 years	>20 years

Figure 33. Benefit Evaluation Metrics for ‘Only FSC 5962’ Option

Evaluation of Alternatives – All 5 FSCs (5961, 5962, 5963, 5998, and 5999)
This scenario calculates the proposed benefit of utilizing blockchain to facilitate the Counterfeit Detection and Avoidance Program with an increased workload to include additional at-risk component classes. These additional component classes include other electronic classes to include 5961, 5962, 5963, 5998, and 5999. As these calculations will show, the increased workload and throughput may prove a net benefit across multiple complexity levels based on the net present value analysis and benefit-cost analysis.

Calculations of Benefits:

- **Cost-Avoided of Labor Associated to Re-procurement**

Labor/Re-procurement Cost-Avoided = (Baseline – Blockchain) = $289,296- $0 = **$289,296**
Present Value = $289,296 x (P/A,7%,10 years) = **$2,031,894**

- **Cost-Reduction of Labor Associated to Error Handling**

Labor/Error Handling Cost-Reduction = (Baseline – Blockchain) = $272,068 - $35,798 = **$236,270**
Present Value = $236,270 x (P/A,7%,10 years) = **$1,659,461**

- **Cost-Reduction of Direct Labor**

Labor/Direct Labor Cost-Reduction = (Baseline – Blockchain) = $470,177 - $213,001 = **$257,176**
Present Value = $257,176 x (P/A,7%,10 years) = **$1,806,298**

- **Present Value of Total Benefits**

PV of Total Benefits = $2,031,894 + $1,659,461 + $1,806,298 = **$5,497,653**

Calculations of Costs Associated with Blockchain Solution:
Implementation Costs are already assumed to be done at present time; thus, no calculations are required other than present value adjustment.

- **Present Value of Tech Support Labor Costs**

  PV of Tech Support Team Costs = $95,846 x (P/A, 7%, 10 years) = $**673,185**

- **Present Value of Cloud Hosting Cost**

  PV of Cloud Hosting Cost = $24,000 x (P/A, 7%, 10 years) = **$168,566**

- **Present Value of Total Costs for Low Complexity**

  PV of Total Costs = $673,185 + $168,566 + $2,500,000 = **$3,341,751**

- **Present Value of Total Costs for Medium Complexity**

  PV of Total Costs = $673,185 + $168,566 + $3,500,000 = **$4,341,751**

- **Present Value of Total Costs for High Complexity**

  PV of Total Costs = $673,185 + $168,566 + $4,500,000 = **$5,341,751**

Low Complexity Analysis:

- **NPV** = Total Benefits at Present Value – Total Costs at Present Value = $5,497,653 - $3,341,751 = **$2,155,902**
- **Benefit-Cost Ratio** = Total Benefits at Present Value / Total Costs at Present Value = $5,497,653 / $3,341,751 = **1.65**
- **Payback Period** = Implementation Cost / Net Annual Cash Flow = Implementation Cost / abs(Annual Benefits – Annual Costs) = $2,500,000 / abs(($289,296 + $236,270 + $257,176) – ($95,846 + $24,000)) = $2,500,000 / $782,742 = **3.8 years**

Medium Complexity Analysis:

- **NPV** = Total Benefits at Present Value – Total Costs at Present Value = $5,497,653 - $4,341,751 = **$455,911**
- **Benefit-Cost Ratio** = Total Benefits at Present Value / Total Costs at Present Value = $5,497,653 / $4,341,751 = **1.11**
- **Payback Period** = Implementation Cost / Net Annual Cash Flow = Implementation Cost / abs(Annual Benefits – Annual Costs) = $3,500,000 / abs(($289,296 + $236,270 + $257,176) – ($95,846 + $24,000)) = $3,500,000 / $782,742 = **6.2 years**

High Complexity Analysis:

- **NPV** = Total Benefits at Present Value – Total Costs at Present Value = $5,497,653 - $5,341,751 = **$143,328**
- **Benefit-Cost Ratio** = Total Benefits at Present Value / Total Costs at Present Value = $5,497,653 / $5,341,751 = **1.03**
- **Payback Period** = Implementation Cost / Net Annual Cash Flow = Implementation Cost / abs(Annual Benefits – Annual Costs) = $4,500,000 / abs(($289,296 + $236,270 + $257,176) – ($95,846 + $24,000)) = $4,500,000 / $782,742 = **6.8 years**

Summary of Evaluation - All 5 FSCs (5961, 5962, 5963, 5998, and 5999)
Based on an analysis of the potential costs associated with a blockchain implementation, it would be feasible to implement a solution if it includes processing of awards for FSCs 5961, 5962, 5962, 5998, and 5999. The greater volume of procurements (roughly 11,500 annually) provides enough benefit for a positive Benefit-Cost Ratio and Net Present Value for all options. Additionally, the payback period is less than 7 years for all options.

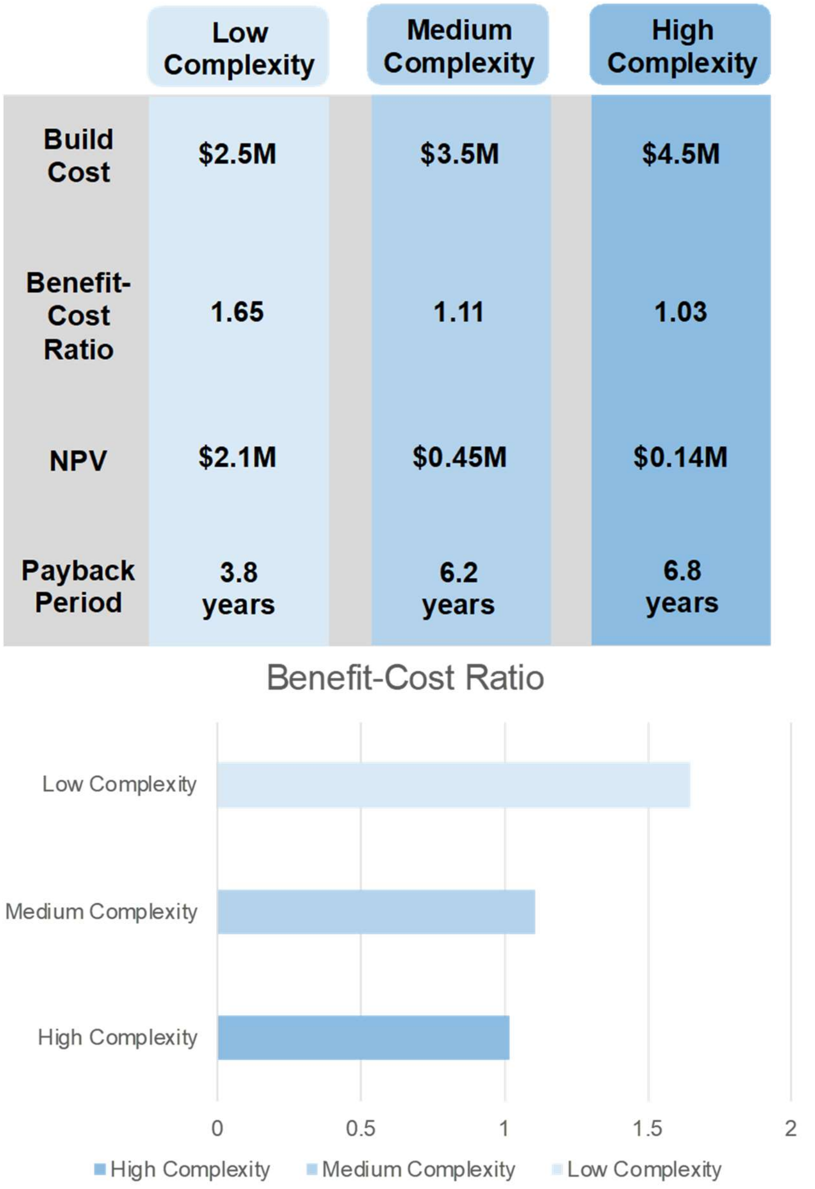

	Low Complexity	Medium Complexity	High Complexity
Build Cost	$2.5M	$3.5M	$4.5M
Benefit-Cost Ratio	1.65	1.11	1.03
NPV	$2.1M	$0.45M	$0.14M
Payback Period	3.8 years	6.2 years	6.8 years

Figure 34. Benefit Evaluation Metrics for 'All 5 FSC' Option

10.0 Transition Plan

Before transitioning a new capability such as the Blockchain Traceability for CDAP application to live production use, it should first go through several phases of development and testing to ensure proper functionality. In general, most projects follow a four-step process from (1) Proof of Concept (or Feasibility Study), (2) Prototype, (3) Pilot, and then (4) Production. The primary goals of each development step are different and the general concepts for each step are described in the figure below.

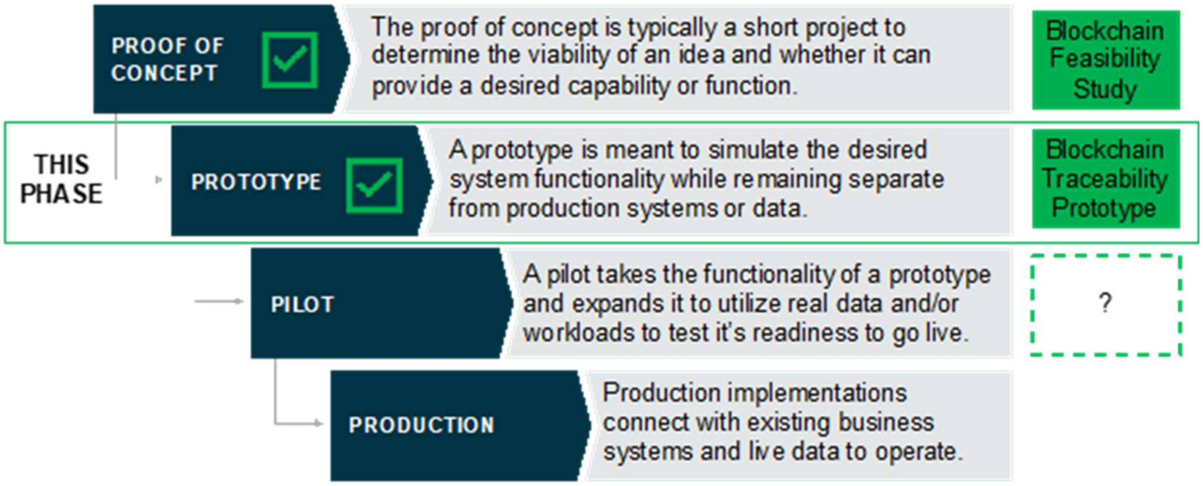

PROOF OF CONCEPT

The proof of concept is typically a short project to determine the viability of an idea and whether it can provide a desired capability or function.

Blockchain Feasibility Study

THIS PHASE

PROTOTYPE

A prototype is meant to simulate the desired system functionality while remaining separate from production systems or data.

Blockchain Traceability Prototype

PILOT

A pilot takes the functionality of a prototype and expands it to utilize real data and/or workloads to test it's readiness to go live.

?

PRODUCTION

Production implementations connect with existing business systems and live data to operate.

Figure 35. Project Development Steps from Proof of Concept to Production

10.1 Technical Components

As a project progresses, the amount of development increases and portions of the previous work may be re-used if possible. During the **Proof of Concept/Feasibility Study** for this project, CDAP processes were analyzed, and during the **Prototype Application** development, scalable microservices were built and deployed. The figure below depicts the types of re-usable assets that were developed during the **Feasibility Study and Prototype** phases of this project and describe the types of activities that should be completed if/when this prototype moves forward into Pilot and Production phases.

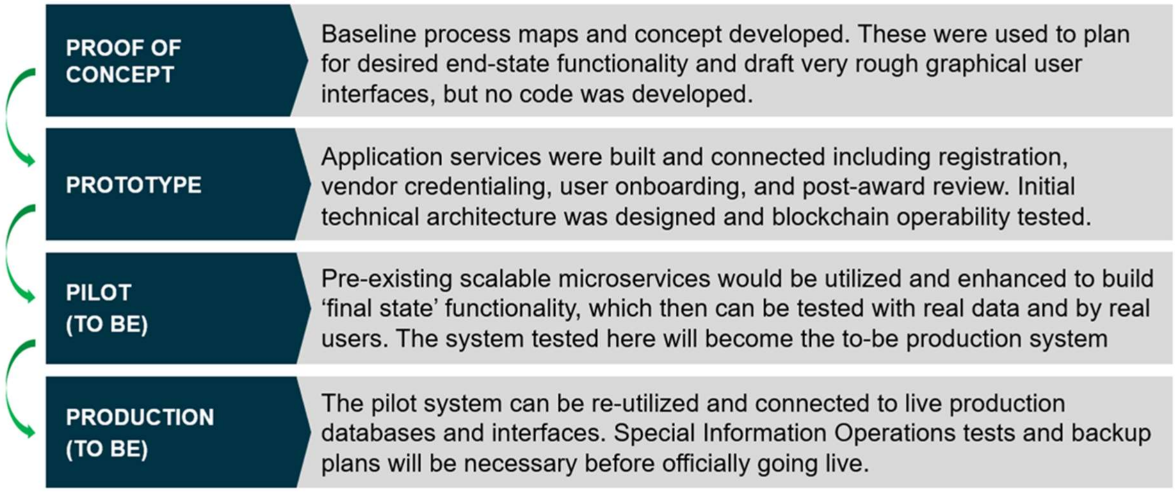

PROOF OF CONCEPT

Baseline process maps and concept developed. These were used to plan for desired end-state functionality and draft very rough graphical user interfaces, but no code was developed.

PROTOTYPE

Application services were built and connected including registration, vendor credentialing, user onboarding, and post-award review. Initial technical architecture was designed and blockchain operability tested.

PILOT (TO BE)

Pre-existing scalable microservices would be utilized and enhanced to build 'final state' functionality, which then can be tested with real data and by real users. The system tested here will become the to-be production system

PRODUCTION (TO BE)

The pilot system can be re-utilized and connected to live production databases and interfaces. Special Information Operations tests and backup plans will be necessary before officially going live.

Figure 36. Reusable Assets and Activities for Each Project Phase

As this project has moved along the development lifecycle (and continues to do so), there are major activities to accomplish at each level. Some activities may occur more than once or appear in multiple different phases – for example, cybersecurity should be analyzed and documented throughout each phase of the project. The figure below lists some of the key activities associated with each phase of the project and indicates a relative number of stakeholder users involved in each phase is to the right of the diagram.

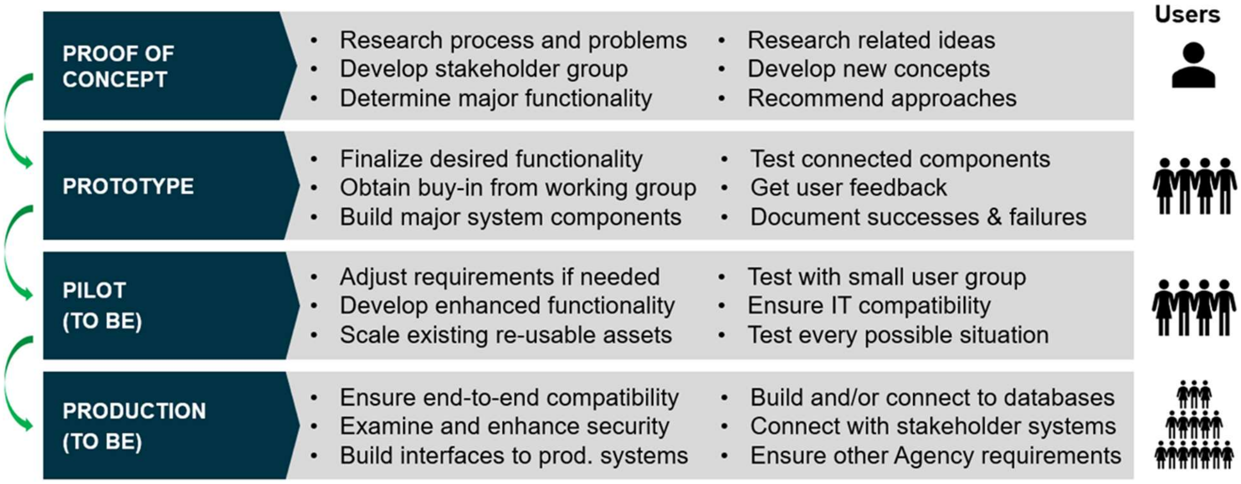

Users

PROOF OF CONCEPT
- Research process and problems
- Develop stakeholder group
- Determine major functionality
- Research related ideas
- Develop new concepts
- Recommend approaches

PROTOTYPE
- Finalize desired functionality
- Obtain buy-in from working group
- Build major system components
- Test connected components
- Get user feedback
- Document successes & failures

PILOT (TO BE)
- Adjust requirements if needed
- Develop enhanced functionality
- Scale existing re-usable assets
- Test with small user group
- Ensure IT compatibility
- Test every possible situation

PRODUCTION (TO BE)
- Ensure end-to-end compatibility
- Examine and enhance security
- Build interfaces to prod. systems
- Build and/or connect to databases
- Connect with stakeholder systems
- Ensure other Agency requirements

Figure 37. Key Activities for Each Project Phase

10.2 Functional Components

The ‘functional transition’ describes how end-users interact with the system along the course of evaluating and developing it for final production use. During initial phases of the project, ideas are developed and tested with input from a small select group of decision makers, before expanding to a larger stakeholder group. The figure below describes some of the core high-level functional activities performed during the previous phases of this project and expected activities for subsequent phases. The relative size of the stakeholder groups is depicted in the figure next to each phase.

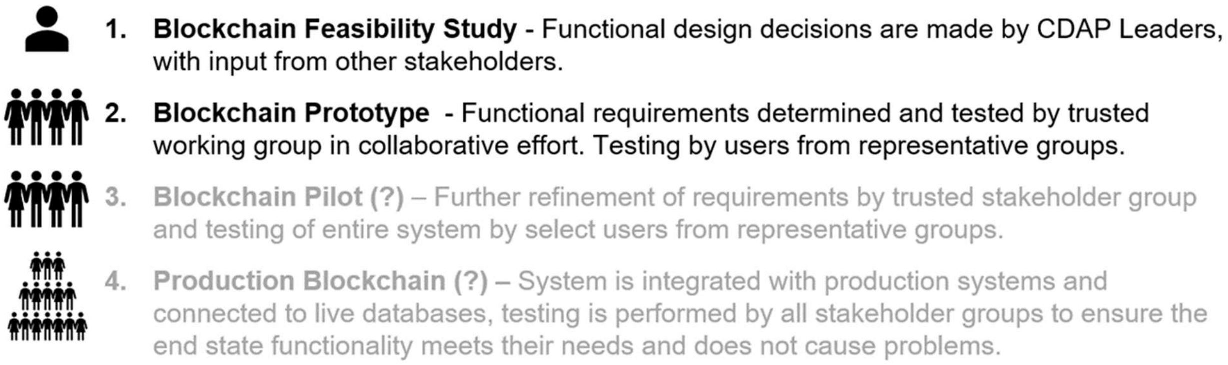

1. **Blockchain Feasibility Study** - Functional design decisions are made by CDAP Leaders, with input from other stakeholders.
2. **Blockchain Prototype** - Functional requirements determined and tested by trusted working group in collaborative effort. Testing by users from representative groups.
3. **Blockchain Pilot (?)** – Further refinement of requirements by trusted stakeholder group and testing of entire system by select users from representative groups.
4. **Production Blockchain (?)** – System is integrated with production systems and connected to live databases, testing is performed by all stakeholder groups to ensure the end state functionality meets their needs and does not cause problems.

Figure 38. Functional Activities Completed During Each Project Phase

In general, proof of concepts test a small portion of the overall solution and do not involve direct end user testing. Prototype phases include a bit more end user involvement and the development of a more complete (but still unfinished and unrefined) future state solution. This phase is where actual end users can begin to test unrefined portions of the capability, and where users were able to test and reflect on the prototype application built during this project phase. The pilot phase will provide an opportunity to test a similar, more refined capability with the same end-users and production-level data, but still would not be utilized to conduct real transactions. It is not until the production implementation that end-to-end integration, testing, and configuration can be completed to allow day-to-day processes and real transactions to occur.

The ultimate goal of the Blockchain Traceability for CDAP application is to provide a new platform for vendors and CDAP users to perform their day-to-day functions. The functional activities of CDAP end-users should be slowly and deliberately increased over the lifecycle of development to ensure appropriate functionality and avoid errors that may setback the critical production processes. Any impacted users (internal or external) should be given ample notice of the new system go-live dates. A high-level description of the functional day-to-day impacts at each project phase are described in the figure below.

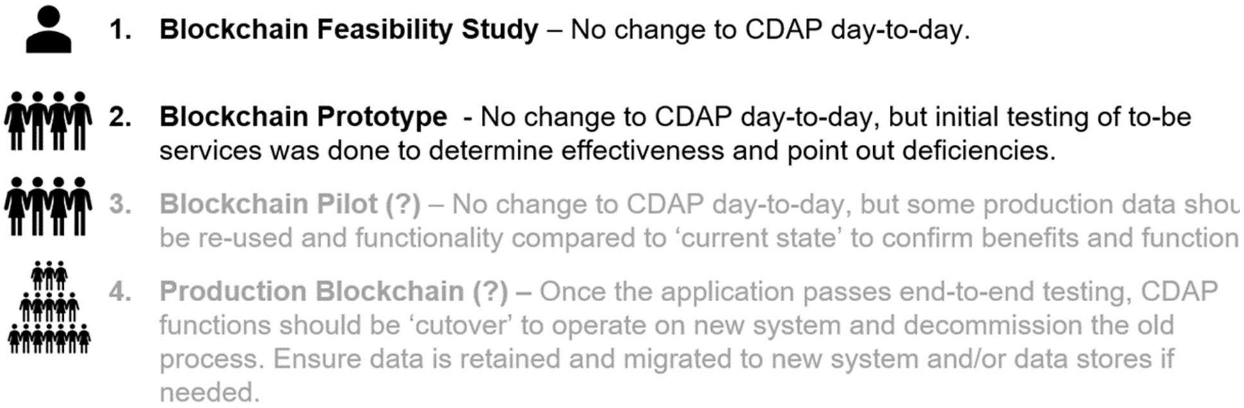

1. **Blockchain Feasibility Study** – No change to CDAP day-to-day.
2. **Blockchain Prototype** - No change to CDAP day-to-day, but initial testing of to-be services was done to determine effectiveness and point out deficiencies.
3. **Blockchain Pilot (?)** – No change to CDAP day-to-day, but some production data shou be re-used and functionality compared to 'current state' to confirm benefits and function
4. **Production Blockchain (?) –** Once the application passes end-to-end testing, CDAP functions should be 'cutover' to operate on new system and decommission the old process. Ensure data is retained and migrated to new system and/or data stores if needed.

Figure 39. Functional Impact on Day-to-Day CDAP Processes

10.3 Transition Phases and Activities

Between each major development phase of a new capability, there is a transition phase where the findings, recommendations, and functional/technical assets (if applicable) should be carried forward and used to inform the next steps of the project. In many cases, the transition phase is a decision point to determine whether there is enough novel development and business benefit to continue with the development lifecycle. In the following sections, we will describe the relevant activities that have been performed so far for this capability and the key functional & technical activities at each step. Additionally, we will describe similar required activities and assets that will be required if moving forward to a Pilot phase or ultimately Production.

10.3.1 Proof of Concept/Feasibility Study Transition to Prototype

To transition the Blockchain Traceability for the CDAP Team application from a proof-of-concept (feasibility study) to a prototype we examined whether blockchain technology was a good choice to enhance the CDAP post-award operations in comparison to pre-existing technology. In this study, our goal was to evaluate blockchain technology in general, its benefits and constraints and how it could impact the CDAP post-award operation, functionally and technically. Furthermore, the study was also used to analyzed potential use-cases applicable to the CDAP operation itself, as well as other areas within the Defense Logistics Agency supply chain operations.

In the Feasibility Study, we identified multiple pain-points throughout the current CDAP process where Blockchain technology could serve as a potential solution. Among these pain-points, the need for a better platform to share information was evidently the most critical area for improvement and aligned perfectly with one of the main pillars to foster the use of Blockchain technology: a shared view of collaborative data. Additionally, we discovered that Blockchain is an effective mechanism to implement a collaborative multi-party ecosystem for sharing information securely, as the technology itself provides a higher level of data integrity across disparate entities through the use of cryptographic and consensus algorithm to keep information synchronized and prevent unauthorized changes to the underlying data.

We analyzed multiple use-cases and ecosystem structures to understand how it will benefit the CDAP team. Moreover, we introduced a decisional framework to evaluate the viability of using Blockchain technology and concluded that, even with a minimal setup that only considered internal DLA stakeholders, blockchain could bring technological advancements in the areas of information sharing and collaboration, decreasing process complexity, reducing the risk of disparate data, enabling the automation of tasks and increasing the review capacity of CDAP operators.

The Blockchain Feasibility Study served as a starting point for introducing this technology to DLA for the purpose of securing one of the most at-risk component classes across the entire enterprise. Furthermore, even with a basic implementation, blockchain technology showed the potential to enable tangible benefits that DLA could realize in a relatively short period of time, in addition of positioning the enterprise at the forefront of blockchain adoption, preparing the organization to quickly leverage all the advantages that this technology will bring in the future.

Key Functional & Technical Activities and Artifacts Considered When Transitioning:

- Researching processes and problems
- Developing Trusted Working Group / Stakeholder Group
- Determining Major Functionality
- Drafting Initial Graphical User Interface (GUI) Concepts
- Researching and Documenting Related Ideas
- Developing New Concepts
- Recommending Approaches
- Developing Initial Recommended Requirements
- Cybersecurity Impacts
- Enterprise Systems Impacts

10.3.2 Prototype Transition to Pilot

Moving to a prototype phase involved a more detailed analysis of the existing processes for each stakeholder group to refine the requirements to implement blockchain technology into the CDAP processes. This transition aimed for building a solution to demonstrate the capabilities offered by the technology. Several activities where performed to establish baseline assumptions, collect functional requirements, design the technology architecture for the proposed solution and developing the prototype.

During the initial portion of this phase, it was important to collect requirements from all of the stakeholders in the ecosystem to ensure that the future-state solution provides capabilities that will benefit everyone. That way, stakeholders will have sufficient incentives to transition to operate under this new approach. Several meetings were scheduled with the different actors in the ecosystem, including OEMs, distributors and Applied DNA Science, to gather feedback and identify functional requirements.

Next, Design Sprints were coordinated and executed to analyze requirements and develop User Personas and User Stories. This provided a more human-centric approach to the design of the prototype's user interface and ensured that the solution not only provides required mechanism to run the CDAP post-award processes, but also to identify conflicting requirements and resolve them. During this step it was also important to kick off an evaluation of all the candidate blockchain platforms. For this project, an evaluation framework was designed to score a total of six blockchain platforms that were good options for the prototype. The framework systemically analyzed the different features, constraints, and nuances of each platform based on ten essential criteria.

Once the requirements were agreed upon by all stakeholder groups, a team of developers and specialists was assembled to construct all the prototype application components. In cases where a development periods are relatively short, it may be a priority to recruit resources with vast knowledge within application development, something often referred to as experience in 'Full-Stack development'. Additionally, rapid prototyping techniques and development frameworks were used along with cloud technology as accelerators to avoid spending a lot of time setting up infrastructure or running administrative configurations.

Following the initial configuration of technical networks, Agile Methodologies were used to organize user stories into Epics and Epics into Sprints. Each Sprint was roughly allocated two weeks to align with bi-weekly working group meetings and present the in-progress results to stakeholders. This approach allowed the team to develop full functionality while keeping the development velocity in line with the plan. The development team, also known as the Scrum Team, conducted several ceremonies to refine the product backlog, score each story's development effort and identify sprint commitments as the project continued. The Scrum Master was responsible for facilitating all of these ceremonies and ensuring that technical barriers were removed in a timely manner to prevent bottlenecks. Furthermore, detailed Acceptance Criteria for each story (aligned to workflows) were defined and proper unit testing was performed to validate their completion. Lastly, it was important to keep business stakeholders engaged and informed. This allowed the Scrum Team to address any concerns as they came up and to efficiently pivot to other approaches, as required.

Once all of the stories were completed, the required components were fully developed and deployed and a fully working prototype was available to demonstrate. At this point, the requirements were reviewed with all stakeholder groups and feedback on the successes and drawbacks of the final prototype application were documented.

Key Functional & Technical Activities and Artifacts Considered When Transitioning:

- Finalize Desired Functionality
- Obtain buy-in from Trusted Working Group
- Build Major System Components
- Test Connected Components

- Get User Feedback
- Document Successes & Failures
- Cybersecurity Impacts
- Enterprise System Impacts
- Software Approval and Readiness
- End User Business Impact

10.3.3 Pilot Transition to Production

The transition from a prototype phase to a pilot phase can vary drastically depending on the amount of development performed during the prototype phase. During the prototype phase of this project, special attention was paid to re-usability of developed assets and scalability of the deployed services. This means that all development was focused on being able to carry forward the technical codebase and frameworks forward once demonstrated in the prototype to be utilized and refined within the pilot phase. Additionally, meetings were held with information technology subject matter experts and topics such as cybersecurity, current-state technology landscape, and compatibility of software were considered. These activities were done to confirm that at the time of the prototype completion, general agreement with these policies and architectures was present, and that no major conflicts existed.

Due to the re-usable nature of these components and the overall alignment with strategic and technology goals (i.e. utilize approved software and understand cybersecurity risks), these processes should provide re-usable frameworks for both the Pilot and Production phases. Within the pilot phase, the codebase and development should begin to be hosted within DLA's firewalls and sandbox environments to simulate a production implementation and examine where further refinement is required. Similarly, new functional and business capabilities should be developed and tested within this environment to ensure no unforeseen errors occur. Each step of the way cybersecurity, alignment with technical policy, and use of approved software should continue to remain a focus. In the pilot and production phases, much closer alignment and collaboration with J6 Information Operations will be required. Additionally, J6 may help guide many of the technical decisions and help determine technical feasibility and deployment based on the ever-changing state of the DLA enterprise infrastructure.

Key Functional & Technical Activities and Artifacts to Consider When Transitioning:

- Refine Requirements as Needed to Meet Business Goals
- Develop Enhanced Functionality
- Scale Existing Re-Usable Assets
- Test with Small End User Group
- Ensure IT Compatibility with Enterprise Framework
- Ensure end-to-end Compatibility of Systems
- Examine and Enhance Security
- Simulate Interfaces to Production Systems
- Build and/or Connect to Databases
- Ensure other Partner Agency requirements
- Test with Other Agency Partners as Needed

11.0 Conclusion

This project continued the digital modernization and enhancement of the Counterfeit Detection and Avoidance Program (CDAP) started with the ***Digital Traceability*** project and continued with the Blockchain ***Traceability Feasibility Study***. This prototype application ***Blockchain Traceability for the Counterfeit Detection and Avoidance*** significantly advanced that work and culminated in the development of a functioning end-to-end digital capability allowing CDAP and vendors to collaborate closely on the same platform. Additionally, several key milestones were achieved and will help inform the next phases of this project including:

- Recommended Requirements for the to-be Blockchain Traceability for CDAP application were developed and confirmed by the Trusted Working Group.
- Potential blockchain development platforms (software packages) were evaluated and Hyperledger Indy was chosen for our identity-based blockchain development, with consideration of Hyperledger Aries as a complimentary component.
- A functioning end-to-end prototype application was developed to meet all major functionality requirements and to demonstrate additional end-state features.
- The functioning prototype was tested by end users and proved successful interoperability of scalable microservices (i.e. registration, review, etc.)
- Other technical and functional considerations such as cybersecurity, governance, and transition were documented and discussed.

Following the success of this prototype phase, a transition to the Pilot phase would be a logical next step in moving this capability closer to production. As previously discussed, a large focus was put on re-usability of developed prototype assets and documentation of enterprise considerations. These assets will prove valuable in any future phases and provide a great starting point for expanding the functionality, user base, and business innovation provided by the end product. Next steps for transitioning this capability into subsequent phases of development and ultimately into production could require some of the following activities:

- Evaluation by J6 for systems readiness and compatibility with approved software.
- Pilot with real data and production-level workloads to determine usability.
- Collaboration with DISA and/or other government organizations to explore distributed systems approach and cybersecurity.
- Exploration of track and trace blockchain and interoperability to provide end-to-end tracking of items from manufacturers to customers.
- Exploration of other operational areas of interest within DLA & related agencies to find opportunities for collaboration & data sharing.
- Development of best practices for implementing a blockchain capability.

Ultimately the primary stakeholders and end users of this capability will provide the best insights into the future business use of this capability and the benefits that it may provide. The Trusted Working Group representing multiple departments and supply chains within DLA will also play a critical role in ensuring requirements from all impacted areas and stakeholders are accounted for. Careful consideration of the potential costs of implementing such a capability weighed against the production benefits of supply chain risk management, credential/identity management, and error reduction will be primary drivers of a potential production use case.

APPENDIX I. List of Acronyms

AWS - Amazon Web Services
ADA - Americans with Disabilities Act
API - Application Programming Interface
ADNAS - Applied DNA Sciences
BFT - Byzantine Fault Tolerance
CoC/T - Certificate of Conformance and Traceability
CDO - Chief Data Officer
CLB - Classic Load-Balancer
CAGE - Commercial and Government Entity
CAC - Common Access Card
CI/CD - Continuous Integration/Continuous Delivery
CDD - Contract Delivery Date
CLIN - Contract Line Item Number
CDAP - Counterfeit Detection & Avoidance Program
DID - Decentralized Identifier
DDWO - Defense Distribution Warren Ohio
DLA - Defense Logistics Agency
DNA - Deoxyribonucleic acid
DoD - Department of Defense
DCA - Design Control Activity
DT - Digital Traceability
DLT - Distributed Ledger Technology
DNS - Domain Name System
EKS - Elastic Kubernetes Service
ECDSA - Elliptic Curve Digital Signature Algorithm
EBS - Enterprise Business Systems
FedRAMP - Federal Risk and Authorization Management Program
FSC - Federal Supply Class
GAO - Government Accountability Office
GUI - Graphical User Interface
HTTP - Hypertext Transfer Protocol
IT - Information Technology
IoT - Internet of Things
JWT - JSON Web Token
L&M - Land & Maritime
MSC - Major Subordinate Commands
MVE - Minimum Viable Ecosystem
NIIN - National Item Identification Number
NSA - National Security Agency
OCM - Original Component Manufacturer
OEM - Original Equipment Manufacturer
OMB – Office of Management & Budget
PII - Personally Identifiable Information
PAR - Post-Award Request

PKI - Public Key Infrastructure

QML - Qualified Manufacturers List

QPL - Qualified Products List

QSLD - Qualified Supplier List of Distributors

QTSL - Qualified Testing Suppliers List

QR - Quick Response

RM - Records Management

RDS - Relational Database Service

RSA - Rivest-Shamir-Adleman

RBAC - Role-based Access Control

SHA - Secure Hashing Algorithm

S3 - Simple Storage Service

SDK - Software Development Kit

SCRM - Supply Chain Risk Management

SDR - Supply Discrepancy Report

TD - Technical Data

TQ - Technical Quality

TWG - Trusted Working Group

USPS - United States Postal Service

VPC - Virtual Private Cloud

W3C - World Wide Web Consortium

ZKP - Zero Knowledge Proof

APPENDIX II. Detailed Platform Evaluation Results

NOTE: These platform evaluations were completed in early 2020 and it is very likely that modifications and enhancements of these platforms have been made since this publication. Additionally, other previously excluded platforms may have been further developed to now meet minimum criteria and new platforms have likely emerged to serve similar purposes. A new platform evaluation should be conducted for any follow-on blockchain projects.

Hyperledger Fabric

Hyperledger Fabric is an open source enterprise-grade permissioned distributed ledger technology (DLT) platform, designed for use in enterprise contexts, that delivers some key differentiating capabilities over other popular distributed ledger or blockchain platforms. It utilizes the AWS EC2 & Azure Blockchain template for Fabric deployment and the code language is written in Golang with a Java & Node SDK. In terms of data distribution, it comes in two forms – channels and private data collections. Only participants in a channel can see the transactions that occur on the channel. The private data collections occur within a channel where a certain subset of data can be shared only with a subset of participants on the channel. The rest of the participants see a hash of the data elements in their transaction log. Hyperledger Fabric contains smart contracts with functions that contain the business logic to calculate updates to shared assets. Updates to the network are committed by any party on the network in a global broadcast fashion. When analyzing scalability and performance, various empirical research papers created by IBM, academia researchers, and production implementers stated a high level of throughput scalability, as well as optimization capabilities. Channels do not scale well, and all identities must be known to set them up (i.e., privacy loss in string building). Hyperledger Fabric has great modularization, and generation of Swagger API with Loopback. Taking traceability and auditability into consideration, states are distributed specific to channels, while carrying logging capabilities within endorsers and peers. Additionally, in regard to recoverability and availability, there is a copy of a ledger on every peer. Copies on different peers is kept consistent, enabling no single point of failure.

Hyperledger Fabric - Criteria Assessment

1. **Private Blockchain: 4**

Hyperledger Fabric is an open source enterprise-grade permissioned distributed ledger technology (DLT) platform.

2. **Use Case Applicability: 2**

Hyperledger Fabric can support a multitude of use-cases, however, some of them are better suited for Fabric than others. In the context of Track-and-Trace and/or Supply Chain related uses cases, HLF is a great fit due to the way it allows the data to be modeled, the way it organizes participants, how it coordinates transaction endorsements between different organizations, and how it provides features to create logical organizational groups and/or consortiums. For Identity/Credentials use-cases, HLF does not provide the proper tools to perform cryptographic verifications and/or attestations of credential's claims. It does have a Membership Service Provider feature and supports the use of public-private key infrastructure (PKI) with Certificate Authorities, but it is not flexible enough to support verifiable credentials data models and related procedures.

3. **Interoperability: 3**

As of now, Hyperledger Fabric supports interoperability protocols only for payments via Hyperledger Quilt (a production-ready library that implements the Interledger Protocol). For any other complex data exchange task (i.e. asset transfer, asset synchronization, etc.) there is still no production-ready solution for Hyperledger Fabric. In terms of integration with other systems, HLF provides the means to easily integrate to different type of external technology services: Data Services, APIs, IAM, etc.

4. Production-Readiness: 4

Hyperledger Fabric went officially into production on March 2017 and is currently used in multiple production enterprise application, most notable IBM Food Trust.

5. Security Considerations: 4

Hyperledger Fabric requires the deployment of Membership Service Providers that are responsible to onboard new members into the network.

6. Platform Maturity: 4

As of today, Hyperledger Fabric production version is currently 2.0, with total of 12,124 commits, 33 releases and a community of 202 contributors.

7. Development Complexity: 3

High-level programming languages are used by Hyperledger Fabric for writing smart contracts. It also has SDK for application development available in two programming languages (Go and Node.js). To properly build and operate a Fabric network, a lot of different steps and type of nodes must be correctly configured beforehand. It can become a very cumbersome procedure to deploy the network, especially for the first time.

8. Smart Contract Enabled: 4

Unlike Corda, which provides a DSL for testing, Hyperledger Fabric does not provide any specific mechanisms for helping with the testing effort. HLF uses a command-line-based tool for any operation involving network configuration, including upgrade of smart contracts deployed. Chaincode can support a large number of use-cases. Chaincode is specific to HLF, and generally is not portable; However, HLF does provide an EVM. ChainCode that will allow you to run Ethereum Smart Contracts on HLF, enabling a higher level of portability.

9. Future Flexibility: 4

HLF has great flexibility on the number of use-cases that can run due to its modular architecture, pluggable consensus engines and data model construct. For CDAP blockchain strategy, HLF can enable the deployment of part's provenance as well as physical asset tracking, amongst many other relevant supply chain use-cases.

10. Governance & Consensus: 4

In HLF, for a transaction to be valid it most have the required level of endorsement (signatures). Consensus over the ordering of transaction is achieved by a specific type of node/service and is part of the transaction processing. The committing nodes perform a validation of both the endorsement policy and double spending

Hyperledger Sawtooth

Hyperledger Sawtooth is an enterprise blockchain platform for building distributed ledger applications and networks. The design philosophy targets keeping ledgers distributed and making smart contracts safe, particularly for enterprise use. Sawtooth simplifies blockchain application development by separating the core system from the application domain. Application developers can specify the business rules appropriate for their application, using the language of their choice, without needing to know the underlying design of the core system. SDKs are available in multiple languages to streamline creation of new contract languages, including Python, JavaScript, Go, C++, Java, and Rust. Furthermore, a provided REST API simplifies client development by adapting validator communication to standard HTTP/JSON. Hyperledger Sawtooth uses Jenkins and has two deployment methods: Debian packages and Docker images. An unofficial AWS EC2 template is used for deployment (An Amazon Machine Image provided by Intel that supports Hyperledger Sawtooth 1.1.4). In terms of data distribution, there is restriction of unauthorized users, and those who access a nodes functionality need to have proper identification. There is clear separation between the application level and the core system level; however, there is no centralized service that could leak transaction patterns or confidential information, all nodes receive all transactions. The business logic contains smart contract abstraction, which allows application developers to write contract logic in a language of their choice. It uses an advanced parallel scheduler that allows for transactions within the same block to modify the same asset. Additionally, pluggable consensus algorithms are used, which allows consensus to be changed at any point. The platform’s scalability and performance allow parallel transaction execution for added throughput, while at the same time preventing double spending. The PoET SGX consensus mechanism allows for high scalability with Byzantine fault tolerance, but at the cost of potential forks that would need to be resolved. In regard to traceability and auditability, Hyperledger Sawtooth has several mechanisms to restrict and secure access to validator peer nodes. The validator node creates proposed blocks from transactions it receives. These proposed blocks are signed by the validator and transmitted to the peer nodes on the network. When analyzing the recoverability and availability on the Sawtooth platform, a single node type is used which simplifies deployment for both On-Prem and on Cloud installations. This integrates with existing databases by keeping internal relational and key value databases in-synch with the Sawtooth network.

Hyperledger Sawtooth - Criteria Assessment

1. **Private Blockchain: 4**

Hyperledger Sawtooth offers the ability to configure private networks.

2. **Use Case Applicability: 2**

Hyperledger Sawtooth created a sample supply chain traceability application as one of its main demos. Traceability in Sawtooth is more applicable to the physical tracking and monitoring of asset due to its architecture which favors more distributed networks.

3. **Interoperability: 2**

Hyperledger Sawtooth uses RESTful APIs for the communication with external services. Its modular design enables good integration capabilities**.** Hyperledger Sawtooth can run Ethereum

Virtual Machine smart contracts via Sawtooth-Seth (an implementation between Hyperledger Burrow and Hyperledger Sawtooth). Although it is possible to integrate Hyperledger Sawtooth with other DLT platforms using Ethereum, it is not recommended due to the lack of standards and maturity in the area of DLT interoperability.

4. **Production-Readiness: 4**

Hyperledger Sawtooth went to production on May 2017. It has been used in multiple production implementations, most notable ScanTrust's supply chain solution for Cambio Coffee in May 2018.

5. **Security Considerations: 2**

Data segregation between participants is extremely difficult to achieve and might require complex access control protocols.

6. **Platform Maturity: 4**

Hyperledger Sawtooth is an open source project initially contributed by Bitwise.io and incubated by Intel. As of today, Hyperledger Sawtooth's community counts with 76 direct contributors and 7,980 commits.

7. **Development Complexity: 4**

Hyperledger Sawtooth provides software development kits for multiple programming languages, including those for mobile development. Hyperledger Sawtooth can run transaction processors via Web-Assembly (portability)

8. **Smart Contract Enabled: 4**

Hyperledger Sawtooth implements smart contracts as transaction families. Each transaction family consist of a transaction processor service deployed on the network and the rules for data storage and transaction validations. In addition, Hyperledger Sawtooth can execute smart contracts on the Ethereum Virtual Machine (EVM) via Sawtooth-Seth.

9. **Future Flexibility: 2**

Hyperledger Sawtooth might not be able to properly support the CDAP blockchain strategy in the future but is a good candidate platform for the development of solutions targeting the physical tracking aspect of CDAP process.

10. **Governance & Consensus: 4**

Hyperledger Sawtooth supports Practical Byzantine Fault Tolerance (PBFT), Proof of Elapsed Time (PoET) and Raft consensus protocols via a pluggable interface. Public-Permissioned (Consortia) as well as fully private network configuration is supported

Hyperledger Indy

Hyperledger Indy provides a distributed-ledger-based foundation for self-sovereign identity. Indy provides a software ecosystem for private, secure, and powerful identity, and the Indy SDK enables clients for it. In terms of data distribution, Hyperledger Indy uses open-source, distributed ledger technology. These ledgers are a form of database that is provided cooperatively by a pool of participants, instead of by a giant database with a central admin. Data lives redundantly in many places, and it accrues in transactions orchestrated by many machines. Strong, industry-standard cryptography protects it and best practices in key management and cybersecurity pervade its design. The result is a reliable, public source of truth under no single entity's control, robust to system failure, resilient to hacking, and highly immune to subversion by hostile entities. The Indy business logic contains cryptographic accumulators and uses a system of revocation issuance to verify credentials through a series of proofs. In regard to scalability and performance, Indy agents require a host. Identity owners can always be their own host, or they can choose a third-party hosting service, called an agency. However, Agents are not strictly required by Indy architecture, as any client can communicate directly with the ledger. With respect to traceability and auditability, the ledger stores Identity Records that describe a Ledger Entity. Identity Records are public data and include Public Keys, Service Endpoints, Credential Schemas, and Credential Definitions. Every Identity Record is associated with exactly one DID (Decentralized Identifier) that is globally unique and resolvable (via a ledger) without requiring any centralized resolution authority. To maintain privacy, each Identity Owner can own multiple DIDs. With regards to recoverability and availability, the Indy validator nodes operate the Plenum protocol, an implementation of RBFT (Redundant Byzantine Fault Tolerance) consensus. A consensus protocol validates each proposed block and its transactions according to endorsement and consensus policies, reaching consensus on the order and results of executing a transaction. It must interface with and depend upon a smart-contract layer for validating transactions in a block. Different methods of consensus are used in Hyperledger products, therefore, validating and ordering transactions are logically distinct processes and interchangeable between consensus mechanisms.

Hyperledger Indy - Criteria Assessment

1. **Private Blockchain: 4**

Hyperledger Indy provides tools, libraries, and reusable components for providing digital identities rooted on blockchains It can support private blockchain network configurations, but it was purposely built to work as a public-permissioned platform.

2. Use Case Applicability: 4

Purposely built for credentials and identity management rooted on blockchain.

3. Interoperability: 4

Hyperledger Indy was designed to be interoperable across multiple distributed ledger platforms.

4. Production-Readiness: 4

Hyperledger Indy went officially into production on April 2019 and is currently used in various production-ready projects, most notable the Verifiable Organization Network.

5. Security Considerations: 3

Indy uses battle-tested cryptographic algorithms for the generation of digital credentials, the signing of transactions committed to the ledger and user access controls.

6. Platform Maturity: 3

As of now, Hyperledger Indy production version is 1.0 with a total average of approximately 23,000 commits or 7,691 commits per repository (currently has 3 main repositories), and a total community of contributors of approximately 224 or 75 per repository.

7. Development Complexity: 4

The indy-sdk has wrapper functions that enables its use in different programming languages. Additionally, deploying a network of indy-nodes is also made simple using Docker containers.

8. Smart Contract Enabled: 1

HLI does not have a smart contract engine built in.

9. Future Flexibility: 1

Hyperledger Indy was purposely built for identity related use-cases, thus has limited flexibility in terms of other use-cases for CDAP (i.e. track and trace).

10. Governance & Consensus: 4

Hyperledger Indy was initially contributed by the Sovrin Foundation, so a lot of its features were built to follow the Sovrin Governance Framework, which is mainly focused on Self-Sovereign Identity, using the distributed ledger as a public identity registry. However, a private network configuration is possible. An Indy network implements the Redundant Byzantine Fault Tolerance (RBFT) called “Plenum”, a leader-based consensus protocol where a node is selected to determine the order of transactions and communicate them to the rest of the network.

Corda

Corda is an open-source distributed ledger platform (not strictly a blockchain) that focuses on privacy and transaction finality. The data distribution is direct peer-to-peer communication rather than global broadcasts. Access to the network is controlled by a service which dictates onboarding policies and procedures. Corda's business logic to calculate the future state is off the network. Updates can be created and submitted by any party on the network in a peer-to-peer fashion and are submitted to the network as a proposal to change the current state to a future state. In regard to scalability and performance, there is high scalability as the peer to peer nature of the network has risk of choke points due to notaries. Corda has an extendable RPC API, and straight forward Swagger API Generation. The Corda platform can integrate with a number of databases, (as long as they have a JDBC adapter) and is built on a JVM stack with development accomplished with JVM languages (e.g. Java, Kotlin, etc.). Traceability and auditability are built into the platform, with a UTXO model, and consumption and un-consumption state management. Additionally, in terms of recoverability and availability, Corda supports disaster recovery, if built into the system. In the case of a node crashing, the corruption or deletion of a node's files are all non-fatal with the right strategy in place. If the checkpoints can be persisted through Node restarts and upgrades, then transactions can also execute across an arbitrary amount of time as flows can be checkpointed.

Corda - Criteria Assessment

1. **Private Blockchain: 4**

Corda supports the configuration of private blockchain networks

2. **Use Case Applicability: 3**

CDAP blockchain prototype can be built in Corda, although its UTXO data model might not provide the proper interface for verifiable credentials

3. **Interoperability: 2**

Corda can integrate with external services via the use of APIs. Interoperability between ledgers is extremely difficult to achieved due to Corda's unique characteristics and architecture

4. **Production-Readiness: 4**

Corda went to production in 2018. Corda is used in multiple production implementations, most notably in the financial sector.

5. **Security Considerations: 4**

Corda provides good security as transactions are only copied to the ledgers of its stakeholders, adding good data segregation between network participants. In terms of transaction finality, Corda guarantees transactions finality by the way it marks used states as historical and by how the Notary node prevents double-spending.

6. **Platform Maturity: 4**

Its enterprise edition is developed and maintained by fintech R3, but most of the platform's core features are open-sourced. As of today, the open-source community is composed of 153 direct contributors and has a total of 8,189 commits.

7. **Development Complexity: 1**

Corda is built on Java and provides SDKs only for JVM compatible programming languages (Java and Kotlin)

8. **Smart Contract Enabled: 3**

Corda's smart contract interface is used for transaction validation purposes. It uses the concept of flows to implement transaction processing protocols

9. **Future Flexibility: 3**

Corda platform is best suited for financial use-cases (i.e. settlements, insurance, payments, etc.). The platform might be able to support both the credentials and traceability use-cases for CDAP, but its configuration and production development might be more complex than other platforms.

10. **Governance & Consensus: 4**

Consensus in Corda is represented in two main aspects: Uniqueness and Validity. Transaction Uniqueness is achieved via the Notary node (tasked with preventing double-spending). Transaction Validity is achieved via Corda Contracts

Quorum

Quorum is an open source blockchain platform that combines the innovation of the public Ethereum community with enhancements to support enterprise needs. Quorum is designed to develop and evolve alongside Ethereum. It is a fork of the Go Ethereum client (geth) and is designed to be developed in line with future geth releases. Because it only minimally modifies Ethereum's core, Quorum can incorporate the majority of Ethereum updates quickly and seamlessly. In terms of data distribution, Quorum uses transaction managers to allow access to encrypted transaction data for private transactions and to manage communication with other transaction managers. Actors have the same view on public states and different views on private states. The business logic uses a majority voting protocol called QuorumChain as a consensus mechanism where nodes within a network can vote on blocks. Solidity language allows for separation of concerns and development of robust smart contracts (RAFT and Istanbul PBFT). Analyzing scalability and performance, since each node needs to execute the Smart Contract, the scalability is somewhat constrained and is highly dependent on how a particular node and smart contracts are configured. Quorum uses the RAFT or Istanbul BFT consensus algorithms that are better performing than mining. Quorum can deploy across different cloud environments or use Docker for cross-environment integration; while also working with existing tools such as Truffle, MetaMask, Remix, and OpenZeppelin. Network permissions are managed on the node level in regard to traceability and auditability. Smart Contract logic needs to be reviewed to prevent the initiator from placing liabilities on others and also to ensure that no sensitive information is being placed on the ledger. Quorum uses ZCASH as a security layer that anonymizes transactions on the blockchain. When gauging recoverability and availability, all the Quorum nodes share the same set of transactions, both private and public transactions are processed by all network participants. Transactions with non-encrypted payloads are called public transactions and update the Public State. Transactions with the hash of encrypted payload are called private transactions. The encrypted payload is exchange off-chain and encrypted/decrypted via the transaction manager.

Quorum - Criteria Assessment

1. **Private Blockchain: 4**

Quorum is a permissioned version of Ethereum focused on data privacy.

2. **Use Case Applicability: 1**

Quorum blockchain platform is typically used for financial use-cases where decentralization or disintermediation is one of the use-case functional requirements. It does not provide proper capabilities for the issuance and verification of digital verifiable credentials, instead custom configuration might have to be built, which could potentially increase security risks.

3. **Interoperability: 2**

Quorum supports ledger interoperability through the inter-quorum asset transfers and has good portability features for other blockchain platforms that supports EVM-based smart contract execution.

4. Production-Readiness: 4

Quorum is a production ready platform as of October 2016. Quorum is currently used in multiple production-ready projects mostly focused on financial applications.

5. Security Considerations: 2

Security is managed at the individual node level. The node is configured to identify other nodes in the network that have been whitelisted for interactions. Quorum implements two components to handles the privacy and security aspect of its transactions: Transaction Manager and Enclave. The transaction manager is responsible for transaction privacy by storing them and allowing access to encrypted transaction data. The Enclave oversees all the cryptographic operations including symmetric key generation and data encryption/decryption, and works in parallel with the Transaction Manager, serving as a virtual HSM. The above components are provided via two main frameworks: Tessera and Constellation

6. Platform Maturity: 4

Quorum is an open-source project that was introduced to the market by JP Morgan in 2016 and ever since maintained in parallel by JPM internal development team as well as the Ethereum community of developers. As of today, it has a community of 383 direct contributors with a total of 11,342 commits

7. Development Complexity: 3

To support private transactions, besides the standard Quorum node, you must also install Constellation and Tessera. Focused on decentralized networks. Private transactions require that nodes that are privy to the private data are available for the transaction to succeed as it is a prerequisite for the recipients to store the communicated payload. This has an impact on transaction finality, and it must be properly accounted for when designing workflows. Quorum only supports the Ethereum Virtual Machine (EVM) runtime

8. Smart Contract Enabled: 1

Quorum supports smart contracts written for the EVM runtime. Maintainability is poor due to the requirement of compiling smart contracts into byte codes that is then included as part of a contract transaction in Ethereum. Upgradability is poor due to the immutability aspect of the Ethereum blockchain which required the destruction of the previous version and the deployment of the new contract.

9. Future Flexibility: 2

Quorum blockchain platform is flexible enough to support multiple use-cases around supply chain traceability and supplier's credentials, we don't see it well fit for the CDAP blockchain strategy due to its architectural ties with the Ethereum platform. Although is already a production-ready platform and an official fork of the Ethereum code base, it stills has strong dependencies with it, which could raise technical issues and scalability concerns in the future.

10. Governance & Consensus: 4

Quorum uses Raft and Istanbul-BFT consensus protocols which are efficient, high throughput and secure consensus protocols. Raft is only crash fault tolerant whereas Istanbul-BFT supports byzantine fault tolerant. Quorum supports consortium and fully private blockchain governance models

Guardtime

Guardtime introduced the keyless signature infrastructure, a method and a globally distributed network infrastructure for the issuance and verification of KSI signatures. Unlike traditional digital signature approaches, e.g. Public Key Infrastructure (PKI), that depend on asymmetric key cryptography, KSI uses only hash-function cryptography, allowing verification to rely only on the security of hash-functions. Guardtime is not strictly a blockchain platform, as it doesn't use a consensus protocol to maintain data across different nodes. Instead it uses a public and global hash calendar registry for the storage and validation of the KSI signatures. Guardtime is a distributed public record of events; an append-only record of events where each new event is cryptographically linked to the previous. New entries are created using a distributed consensus protocol. In regards to data distribution and business logic, users interacts with the KSI system by submitting a hash-value of the data to be signed into the KSI infrastructure and is then returned a signature which provides cryptographic proof of the time of signature, integrity of the signed data, as well as attribution of origin (i.e. which entity generated the signature). The properties of the signed data can be verified without reliance or need for a trusted authority. Analyzing the Guardtime platform scalability and performance, the KSI signatures can be generated at exabyte-scale. Even if an exabyte (1,000 petabytes) of data is generated around the planet every second, every data record (a trillion records assuming 1MB average size) can be signed using KSI with negligible computational, storage and network overhead. The properties of the signed data can be verified even after that data has crossed geographic or organizational boundaries and service providers. When gauging traceability and auditability, KSI does not ingest any customer data; data never leaves the customer premises. Instead, the system is based on one-way cryptographic hash functions that result in hash values uniquely representing the data but are irreversible such that one cannot start with the hash value and reconstruct the data (data privacy is guaranteed at all times). With respect to recoverability and availability, the number of participants in KSI blockchain distributed consensus protocol is limited. By limiting the number of participants, it becomes possible to achieve consensus synchronously, eliminating the need for Proof of Work and ensuring settlement can occur within one second.

Guardtime - Criteria Assessment

1. Private Blockchain: 2

Guardtime offers the ability to connect its hardware components to create private networks configurations, although its core "distributed ledger" resides on a public domain.

2. Use Case Applicability: 2

Guardtime provides hashing mechanisms that can be used to create unique digital fingerprints which enables great data integrity features, however only hashed are stored on the network, limiting its practical use as a “Shared View of the World”.

3. Interoperability: 4

Guardtime has good interoperability and portability features that allows it to integrate with other blockchain platforms or external services.

4. Production-Readiness: 4

Guardtime is a production-ready platform currently used in several projects, most notably by Estonia’s government for the management of government records.

5. Security Considerations: 4

Guardtime company guarantees that the platform works with 99.99% availability, resistant to denial of services attacks and able to withstand attacks by quantum computers.

6. Platform Maturity: 4

The platform went into production in 2007 and is privately developed and maintained by Guardtime corporation.

7. Development Complexity: 2

The platform is provided only by Guardtime which could potentially create vendor’s lock-in and supportability issues in the future.

8. Smart Contract Enabled: 0

The platform does not directly support any type of smart contract execution.

9. Future Flexibility: 1

The platform could greatly support CDAP blockchain roadmap, but very limited capacity due to the lack of practical data visibility as only a root hash of multiple hash trees is stored on the ledger.

10. Governance & Consensus: 1

The platform requires other technologies to support the implementation of governance structures. Consensus is only required for the validation of the ledger’s root hash.