GS1 Trusted Source of Data

Pilot Report
January 2012

Document Summary

Document Item	Current Value
Document Title	GS1 B2C Trusted Source of Data Pilot Report
Date Last Modified	Jan 2012
Current Document Issue	Issue 1
Status	Draft
Document Description	This document provides the technical details of the GS1 B2C Trusted Source of Data pilot that ran from July – December 2011.

Contributors

Name	Organization
Dipan Anarkat	GS1 GO
Cameron Green	GS1 GO
Joe Horwood	GS1 GO
Malcolm Bowden	GS1 GO

Log of Changes in Issue 1

Issue No.	Date of Change	Changed By	Summary of Change
1	25 January 2012	Dipan Anarkat, Joe Horwood, Cameron Green, Malcolm Bowden	FINAL

Disclaimer

Table of Contents

1. Executive Summary

1.1. Objective

GS1 aims to become the Trusted Source of Data (TSD) to support the communication of authentic product data provided by brand owners to consumers/shoppers, retailers, internet application providers, and government using internet and mobile devices.

In June 2011 a Proof-of-Concept (PoC) established the basic architecture for a system to support this goal. Between July and December 2011 a global pilot was organised to perform a more robust test of this system with more participants and product data.

1.2. Methodology

More than 30 brand-owners in 8 countries provided information on over 900 products. In addition to real-time testing within 5 different mobile applications by the pilot participants, a software model was developed to fully test all system connections for each of the products.

1.3. Results

The pilot showed that it was feasible for multiple internet application providers (IAPs) to access consumer-facing digital product information supplied by multiple brand-owners, using a global and interoperable approach. Performance and quality metrics demonstrated that the system functioned correctly with appropriate response times.

1.4. Learnings & Recommendations

The pilot provided invaluable learnings that should be taken into account for the deployment of the Global Trusted Source of Data (TSD) Framework. Key recommendations are as follows:

Data Provision

- Develop guidance for populating consumer-facing product data and provide data quality requirements for B2C applications
- Provide guidance on how private label products marked with numbers for internal use can be used in the framework

Technical Architecture

- Eliminate the need for a gateway service managed by GS1 by ensuring local data aggregators can directly service Internet Application Provider (IAP) requests
- Implement the Federated Object Naming Service (F-ONS) as the registry function to match the bar code number with the location of the product data
- Develop a standard GS1 solution to ensure bar code numbers can be easily and unambiguously resolved by ONS
- Develop guidelines addressing all aspects of security in the TSD framework
- Consider options to increase system performance by supporting caching of data

Data Usage

- Use a data format that simplifies integration into mobile applications by IAPs
- Better understand B2C requirements related to target market and provide appropriate standards and guidelines

1.5. Acknowledgements

GS1 would like to express its thanks to all those who participated in the pilot and in particular to the B2C Project Board, the GS1 B2C Experts Group and the staff and resources involved with the pilot at GS1 Australia, GS1 Colombia, GS1 Canada, GS1 Germany, GS1 France, GS1 Spain, GS1 US, GS1 UK, 1Sync, checkitmobile, epcSolutions, ipiit, Ken Traub Consulting, Mirasense, Proxima Mobile and Verisign.

1.6. Next Steps

GS1 and relevant stakeholders will now focus on the deployment of the Global B2C TSD Framework by development of global standards and guidelines and a common operational model.

Local aggregation in some cases has already begun and will conform to the global framework as important aspects are deployed.

2. Objective

GS1 aims to become the Trusted Source of Data (TSD), to support the communication of authentic product data provided by brand owners to consumers/shoppers, retailers, internet application providers, and government using internet and mobile devices (phones, laptops, etc).

In June 2011 a Proof-of-Concept (PoC) established the basic architecture for a system to support this goal. Between July and December 2011, a global pilot was organised to perform a more robust test of this system, with more participants and product data.

The system used for the pilot was designed with the following considerations in mind:

- leverage existing standards based on brand data used by the Global Data Synchronization Network (GDSN) and other accredited data sources,
- focus on basic and nutritional information desired by consumers and available by brands in the GDSN.

The pilot was to be deemed successful if two or more mobile applications designed by different Internet Application Providers (IAPs) were able to:

1. send requests for information about two or more specific products from different brand-owners
2. receive this information from different locations and display it as authorised by the brand owner.

3. Methodology

The Global Data Synchronisation Network (GDSN) has proved to be an effective tool for sharing product data in the B2B/supply chain space. GS1 therefore intends to leverage the existing Global GDSN infrastructure.

However, GDSN today it is not sufficient for sharing product data for integration in consumer-facing applications for two main reasons:

- **Reliance on existing trading relationships**. GDSN follows a model where trading partners allow the sharing of product data with other trading partners. Consumer-facing applications require data to be shared with internet application providers who do not have a trading relationship with each brand owner.
- **Need to aggregate data that exists within the GDSN network with data from additional sources**. Digital product information or information beyond B2B master data for consumers may need to be aggregated from various accredited sources including the GDSN and other third parties.

With this in mind, a system was designed where:

1. relevant product information was aggregated from GDSN Data Pools and certified third-party sources
2. a registry function allowed each product's Global Trade Item Number (GTIN) to be matched to the location of the aggregated product information

Below is a basic diagram of the architecture used:

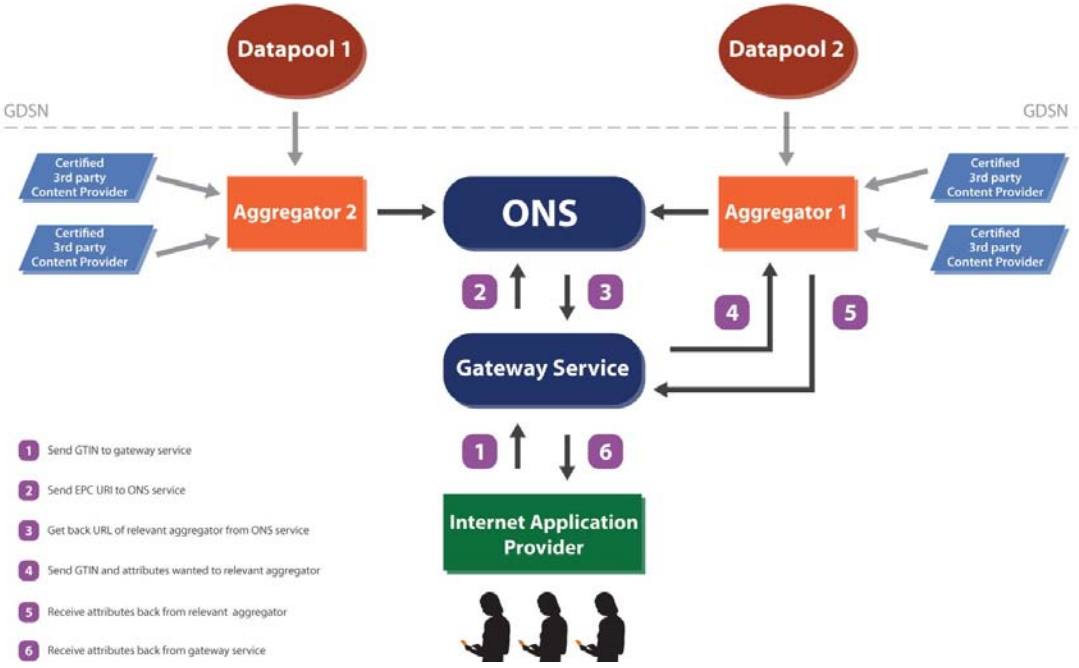

Figure 1: Pilot Architecture and Data Flow

A simple way to describe the data flow is to imagine a consumer is scanning a bar code with the Internet Application Provider's (IAP)'s application. The data flow then proceeds as follows:

- Steps 1 and 2: The product identifier (GTIN) is decoded from the bar code and passed to the Object Naming Service (ONS).
- Step 3: The ONS responds with the location of the information in the relevant Data Aggregator.
- Step 4 and 5: The GTIN is then passed with the request for information to the correct Data Aggregator and is returned a packet of the information for the requested product.
- Step 6: This information is then rendered by the Internet Application Provider (IAP) in the consumer-facing mobile application.

3.1. Pilot Components

To efficiently deliver the pilot, existing systems and implementers were selected to pass the information from source to consumer. More than 30 brand-owners in 8 countries provided information on over 900 products. In addition to real-time testing within 5 different mobile applications by the pilot participants, a software model was developed to fully test all system connections for each of the products.

Below is a list of the participants within each specific pilot role, including the required actions to enable the pilot.

1. **Brand-owners**
   Agua Mineral San Benedetto, Almacenes Éxito, Bonduelle, Capsa, Cadbury, Carrefour, Casino, Coca-Cola, Danone, Emmi, Eroski, Fromageries des Chaumes, Fromageries Papillon, Fromageries Perreault, Les Fromagers de Thierarche, Groupe BEL,Heineken, Hormel Foods, Lactalis, Lesieur, Kellogg's, Kraft, Marie Morin, Nestlé, Noel, PepsiCo, Premier Foods, Rausch, R&R Ice Cream, Smucker's, Taeq, Team, Tesco and Unilever
   Action: Authorise basic product data and nutritional attributes and populate into the framework via GDSN Datapool, Certified 3rd party or Data Aggregator to be accessible to IAPs.
2. **GDSN Datapool**
   1SYNC **Action:** Send basic product data and nutritional attributes to aggregator.
3. **Certified 3rd Party Content Provider**
   Brandbank
   Action: Authorise and send product image for specific products to aggregator.
4. **Data Aggregators**
   1SYNC, GS1 NutriFacts Canada, MO2O, CABASnet and ProductOnLine.
   Action: Combine brand authorized data from GDSN Datapool with additional 3rd party data (such as product image). Load ONS with GTINs for products with information in aggregator.
5. **ONS**
   Peer ONS Root name servers, GS1 Test ONS and EU Root ONS, managed by GS1 Global Office and GS1 France respectively (until Federated ONS is developed and implemented)
   Action: Receive GTINs with location information (which aggregator), return location information to IAP requests.
6. **Gateway Service**
   GS1 B2C Alliance Sandbox used since most of the work, including ONS connectivity, was already done
   Action: Act as the "Gateway" in the pilot to provide the web service for IAPs to connect to TSD framework. Exchange product information requests between IAP and ONS. Sends GTIN requests to ONS, receives location of product information (which aggregator), send request to aggregator, receive product data, return product data to IAP.
7. **Internet Application Providers (IAPs)**
   Checkitmobile, epcSolutions, ipiit, Mirasense and Proxima Mobile
   Action: Scan bar code on product to retrieve GTIN, send GTIN to Gateway Service, receive and render product data for consumer on mobile application.

3.2. Technical Architecture

3.2.1. Overall Approach

The goal was to allow a mobile application to scan an EAN/UPC bar code and receive trusted information for that GTIN as provided by a data provider. Each GTIN has a pointer to the data provider for that GTIN registered in ONS (a many-to-one relationship, as one data provider may serve data for many different GTINs).

One possible design is to have the phone itself perform the ONS lookup, and based on the pointer received from ONS contact the appropriate data provider directly. However, this was deemed to be an unrealistic design, for several reasons:

- The application programming environments typically available for mobile software do not usually provide the ability to interface to ONS. Specifically, they generally do not provide a way to do a DNS lookup to obtain NAPTR records.
- Mobile phone applications are invariably designed to interact with a specific application back-end service provided by the mobile application author. The back-end service then mediates any interaction with external data sources.
- Having an intermediate back-end service also provides for usage logging and additional services making the application architecture much more flexible.

For this reason, the pilot architecture had the mobile application send the GTIN to a back-end application. The back-end application performed the ONS lookup, and then queried the appropriate data provider, finally delivering the result to the handset.

The GS1 B2C Alliance Sandbox provided exactly the back-end application functionality that was needed. It:

- was already designed to accept a query in which a GTIN is specified and deliver information pertaining to that GTIN obtained by querying one or more external data sources.
- included a flexible data transformation mechanism, which allowed the pilot to adapt to slightly differing data formats made available by each data provider, and to transform these formats into a common format for delivery to the mobile application.

All that was necessary was to enhance the GS1 B2C Alliance Sandbox with the option to perform an ONS lookup to determine which data provider to use (the prior functionality was to query **all** data providers and aggregate the results).

For the purposes of the pilot, the GS1 B2C Alliance Sandbox was therefore the Gateway Service shown in Figure 1 (page 7).

3.2.2. Bar code / ONS Integration

An ONS query is based on transforming an EPC identifier into a DNS domain name. In terms of GS1 Identification Keys, the ONS lookup requires that the portion of the key corresponding to the GS1 Company Prefix (GCP) be separated from the remainder of the key. An EAN/UPC bar code with a GTIN, however, does not indicate where to make the division between the GCP and the remainder of the key. This presents a challenge in doing an ONS lookup based on a GTIN bar code.

The solution adopted for the pilot was to take an iterative approach. The number of digits in the Global Company Prefix (GCP) is at least six and at most eleven. For any given GTIN, therefore, there are six different possible division points, corresponding to a GCP length of 6, 7, 8, etc. The ONS lookup software implemented for the GS1 B2C Alliance Sandbox tried each of these possibilities in turn, until the lookup succeeded. In the pilot, the actual GCP lengths were either 7 or 8, so in the pilot each GTIN lookup required two or three trials.

Despite the extra lookups, the overall latency of responding to a query was still acceptably low. End-to-end response times from the moment of scan until data was received from the TSD

framework and rendered on the mobile application provided a satisfying consumer experience without disengaging consumer interest.

3.2.3. Software Development Considerations for ONS

To do an ONS lookup, a software application must use the host operating system's DNS resolver to query for DNS NAPTR records. While all operating systems include a DNS resolver, looking up NAPTR records is a comparatively unusual operation, compared the usual DNS queries for A records (IP address lookup), MX records (mail server lookup), CNAME records (hostname alias lookup), and NS records (nameserver lookup).

Initially the Gateway Service was unable to performing this kind of ONS lookup since it used the Google Application Engine for Java, which is a Platform-as-a-Service (PaaS) cloud-based platform for rapid application development. This environment provides a restricted set of networking primitives, which gives the platform greater flexibility to perform dynamic load management.

One possible solution was to port the Gateway Service to a less restrictive platform. For expediency, however, this approach was not taken. Instead, only the ONS lookup portion was implemented on a different platform, and this portion exposed a REST-style Web Service for use by the Gateway Service proper. Since the Google platform permitted access to Web Services (as virtually all development platforms do), this made it possible to perform ONS lookup with minimal software development effort.

The ONS lookup portion was implemented in Java running on the Amazon Elastic Computing Cloud (EC2), using the Elastic Beanstalk platform.

3.2.4. ONS Setup and Resolution

One of the architectural considerations for the B2C pilot was the need to support multiple and globally distributed peer ONS instances in the B2C pilot, each having a DNS top-level domain (TLD), and with GTIN's scattered across them. This need is predominantly driven by the fact that some advanced GS1 MOs had existing national or regional ONS instances (for example, GS1 in Europe) which they wanted to leverage for the global B2C pilot. This is also GS1's vision of a Federated ONS (F-ONS) system for the future, to overcome the shortcomings of the existing 'centralized' ONS model.

To support multiple global peer ONS Root name servers over a single root ONS name server and factoring in the need to allow for national/regional management of ONS, a hybrid ONS solution was adopted for the B2C pilot that closely mimicked an F-ONS type of system. This was accomplished by updating the Gateway Service with logic to delegate the ONS query to either the B2C Pilot 'Test' ONS or any other participating national/regional ONS. The selected ONS then further resolved the query. With this setup all ONS instances are considered peers; there was no single root. This setup is closest to the F-ONS vision.

For the ONS-based resolution of GTINs to their corresponding web service URL to function, a typical setup would consist of the multiple globally-distributed peer ONS name servers in the GS1 federation. Each peer ONS root name server would then link to multiple local ONS name servers, one for each company prefix owner. The setup of the ONS would be as follows:

Peer ONS Root	Local ONS
Contains "NS" (Name Server) DNS records	Contains "NAPTR" (Naming Authority PoinTeR) DNS records
One NS record for each GS1 Company Prefix	One NAPTR record for each *GTIN*
NS record has the internet address of the Local ONS name server	NAPTR record has the *Service-ID* and *URL*
Name Server maintained by GS1's ONS service provider	Name Server maintained by Brand Owner's ONS service provider
GS1 has authoritative control over the ONS root domain	Brand Owner has authoritative control over Local ONS domain.

The “typical” setup described above would require the brands participating in the pilot to setup a local ONS name server within their enterprise network (typically the enterprise DNS). To speed up pilot implementation and reduce brand owner investment, the peer ONS root name servers in the pilot were set up to provide terminal ONS entries. With this approach, the peer ONS root hosted the DNS zone record on behalf of the Brand Owner rather than requiring delegation to the brand owner’s Local ONS name server. To minimize impact to the existing production Root ONS system (onsepc.com), a peer ONS root name server having the top level domain (TLD) “onstest.com” was setup specifically for the purpose of the B2C pilot. Participating brand owners were then registered with either of the following two Peer ONS Root name servers.

Peer ONS Root	Managed by	TLD	DNS Operator	Product Registrations	Country
GS1 Test ONS	GS1 Global Office	onstest.com	Verisign	796	US, UK, Canada, Colombia, Germany, Spain & Australia
GS1 EU ONS	GS1 France	onsepc1.eu	Orange	114	France

The DNS operators for the two peer ONS root name servers above provided a managed DNS service for ONS registrations.

- Setup of DNS Resource Records was managed using a web-based interface.
- ONS resolution services provided with 100% uptime and near real-time updates to domain data.

3.2.5. Mobile Applications

The pilot allowed multiple Internet Application Providers to connect their mobile applications to the system. Each application retrieved data via the REST interface from the Gateway Service and rendered the XML formatted data in a native Android or iPhone user interface. To access the Gateway Service the mobile application used an application key that was manually assigned a-priori. The Gateway Service returned XML data formatted as per the B2C TSD pilot framework specification. The returned data was presented to the consumer in a format similar to the standard nutritional panel that appears on products.

Example screenshots from two of the mobile applications are shown below.

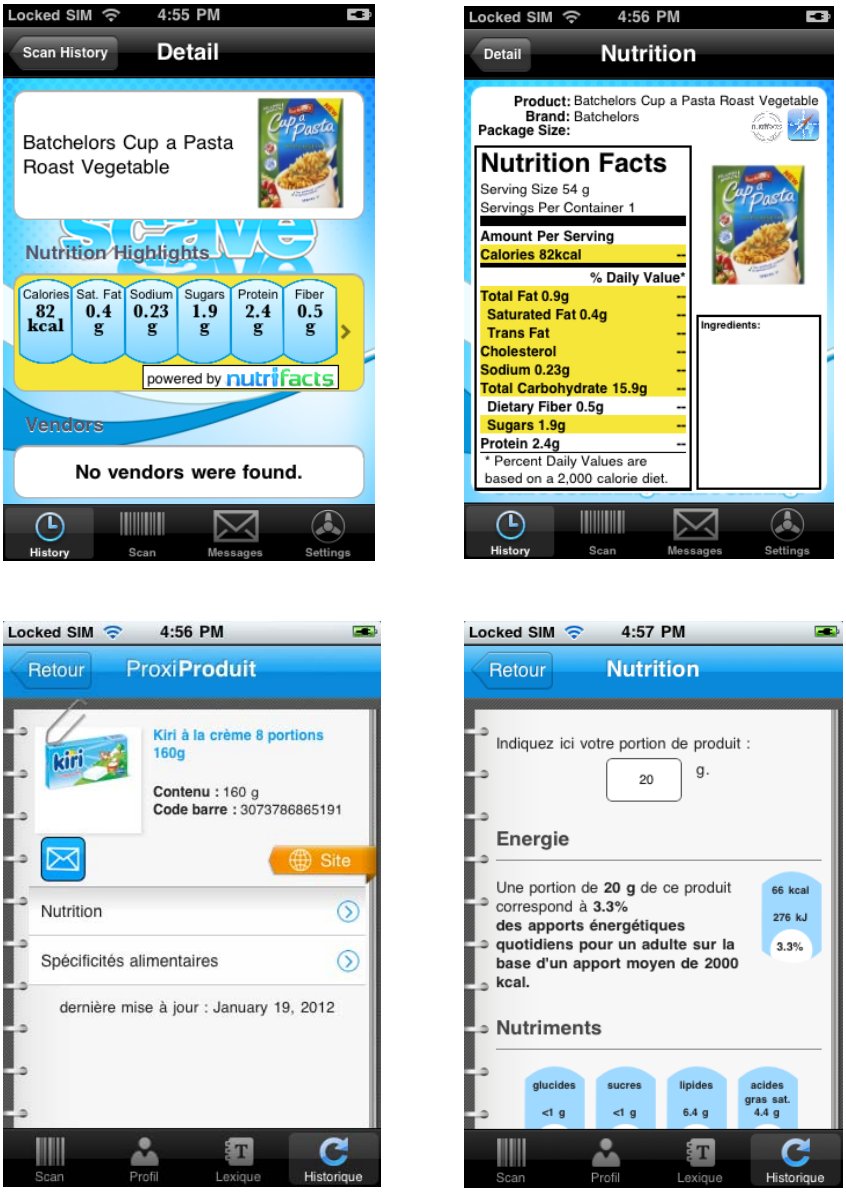

Figure 2: Example screenshots from two mobile applications

3.2.6. System Interconnections

In order to test the pilot system architecture and validate the framework, the results needed to be provided in real time by the live system components. Stakeholders participating in the pilot were interconnected using well-defined web protocols and/or agreed web service interfaces. The table below provides an overview of the connections used in the pilot and the input and output results produced.

Connection	Interface / Protocol	Request Data Format	Request Data Content	Response Data Format	Response Data Content
IAP to Gateway Service	B2CDataInterface → productByGtin web service API (as per specification)	HTTP REST styled message (HTTP GET parameters)	GTIN	XML	ProductData structure containing B2C product information (as per specification)
Gateway Service to Peer ONS Root	DNS	DNS query	ONS domain name (DNS Fully Qualified Domain Name derived from GTIN)	DNS NAPTR Record	Data Aggregator Web Service URL
Gateway Service to Data Aggregator (same as IAP to Gateway Service)	B2CDataInterface->productByGtin web service API (as per specification)	HTTP REST styled message (HTTP GET parameters)	GTIN	XML	ProductData structure containing B2C product information (as per specification)

Establishing the interface connections described in the table above required systems development and testing of all systems components. Additionally, B2C product data was authorized by the brand owners and sourced via their GDSN Datapool and provided by their chosen data aggregator. The mobile applications scanned bar codes from these products and retrieved data from the trusted source via the Gateway Service.

4. Results

In addition to full system testing by pilot participants, GS1 Global Office conducted end to end testing of all connections for all the products utilised in the pilot. The following general results were observed:

- GTINs scanned from the product were resolved properly via ONS, providing the web service URL of the trusted source Data Aggregator
- The Data Aggregator provided authentic B2C product data coming from GDSN and/or the Brand Owner
- Authentic B2C product data was provided to the IAPs and consequently displayed in the consumer-facing mobile applications.
- Testing showed that the average time for data transfer between internet application providers and data aggregators was 1.4 seconds. This means that the time between the bar code being scanned and the product data being displayed in the mobile application was on average between 3 and 5 seconds under normal operating conditions and so deemed appropriate for consumer applications.

Additionally, GS1 leveraged the mobile emulation capabilities of the Android platform to thoroughly test the B2C system. Software program code was used to simulate an IAP connection to the Gateway Service / GS1 B2C Alliance Sandbox. Real time data retrieved from the pilot framework was provided to an Android test application created for the purpose of the pilot. Data collected from the system was collected and analysed and is shown in the following sections.

4.1. Product Data by Stakeholder

In total, 910 products were included in the pilot. The products were distributed as follows by ONS, country and data aggregator.

Product Registration by ONS

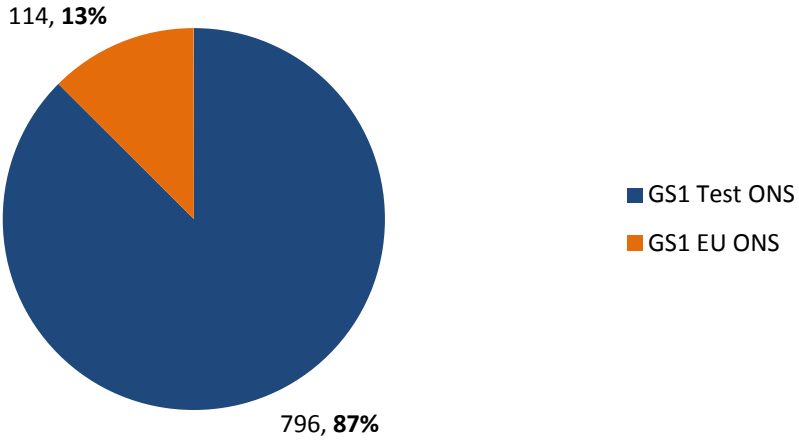

Product Registration by country

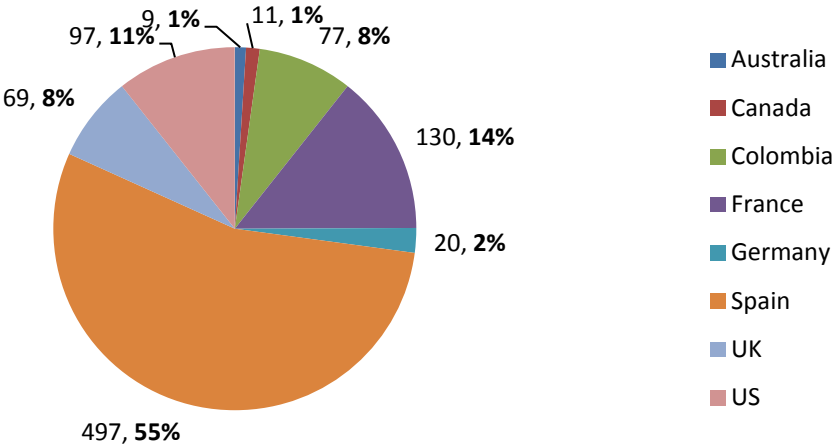

Product Registration by Data Aggregator

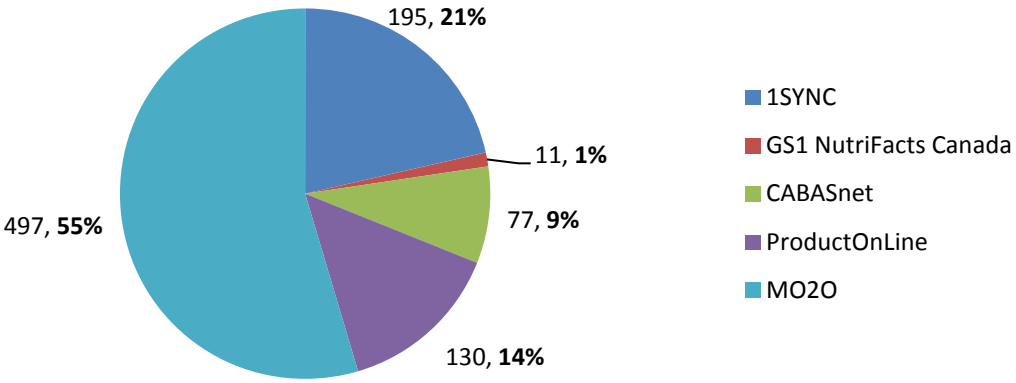

4.2. System Performance

4.2.1. Resolution

Failed resolutions were due to improper resolution of GTINs at the data aggregator web service level, rather than at ONS level. This is likely due to GTINs not being set up properly within the data aggregator's database.

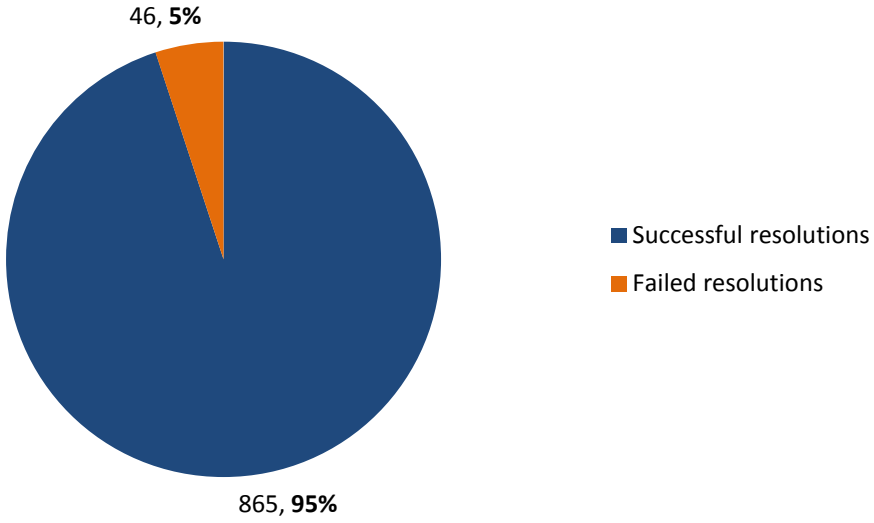

4.2.2. Response Time

Testing showed that the average time for data transfer between internet application providers and data aggregators was 1.4 seconds as shown below.

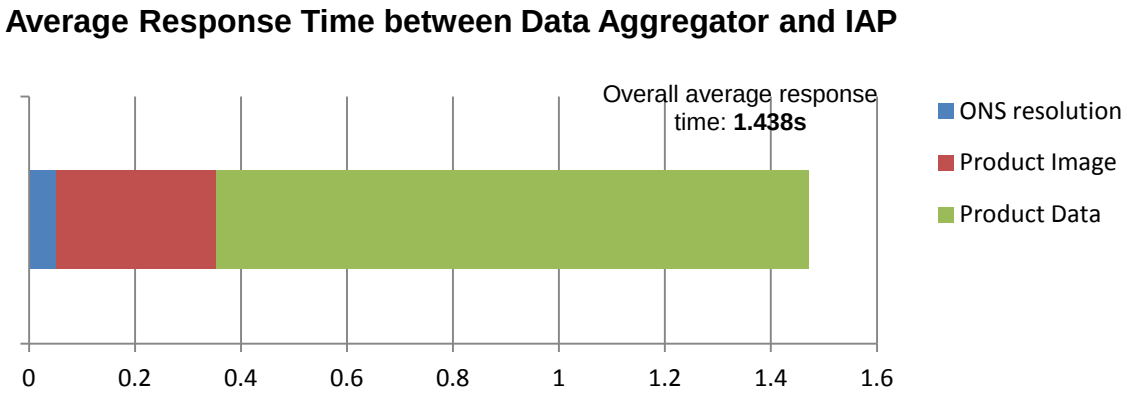

Response times in the real world depend upon a combination of factors, such as network speed and latency, such as server performance, caching, radio signal strength, location, bandwidth, mobile hardware and mobile phone software capabilities. Under normal operating conditions, the time between the bar code being scanned and the product data being displayed in the mobile application was between 3 and 5 seconds on average.

4.3. Data Quality

As already noted, high data quality is a key success factor for this project.

In their “Mobile-Savvy Shopper Report” published by GS1 UK showed the data quality challenge for consumer-facing information by checking product name and image for 375 products in 3 mobile applications. Comparing the results of the UK study with data from the B2C pilot shows that signigicant progress has been made:

- Missing/wrong image: 88% (UK study) / 72% (B2C Pilot)
- Missing product name: 75% (UK study) / 2% (B2C Pilot)

Despite this progress, the fact that data quality was far from perfect for certain attributes shows the need for understanding and collecting data quality requirements for B2C applications (recommended in Section 5).

Basic Product Data

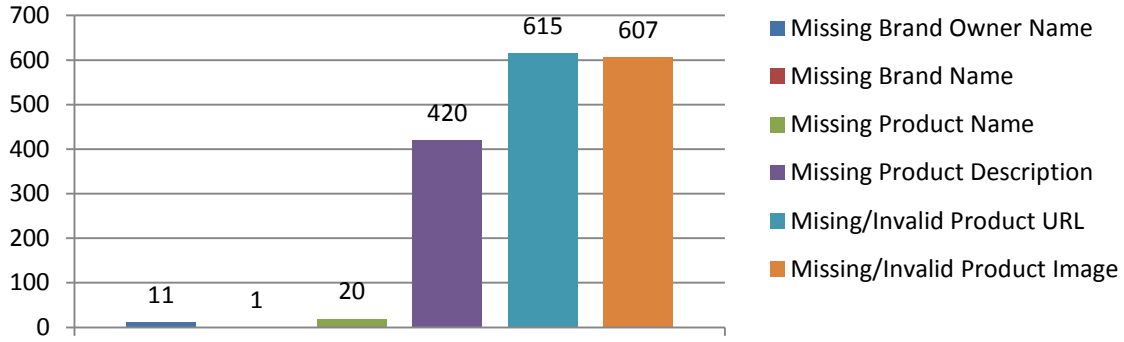

Nutritional Data

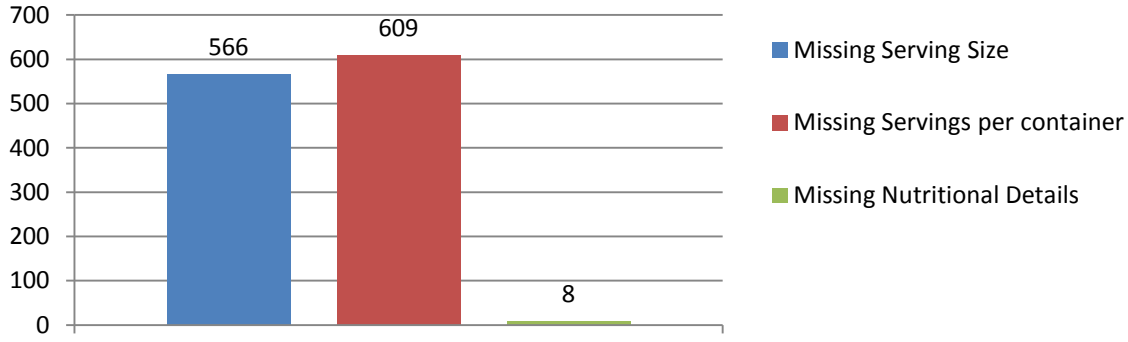

Product URL, product image, servings size and servings per container are attributes that are currently optional within GDSN, hence the large number of products missing this information.

GDSN has multiple attributes for product description, thus creating confusion on which attributes should be used for consumer-facing information and so causing product description to be missing in many cases.

5. Learnings & Recommendations

The pilot was extremely successful and yielded a number of important learnings that will be considered for the Trusted Source of Data (TSD) framework that GS1 intends to put in place. These learnings have been categorized below from the perspective of the individual system components in the TSD framework.

5.1. Data Provision

5.1.1. Brand Owners

The pilot revealed the importance of data quality in GDSN. The following authentic B2C product data attributes from the brand owner's GDSN Datapool is provided to the consumer utilizing the TSD framework.

- Basic Product Data: *GTIN, Product Name, Brand Owner Name, Product Description, Product URL, Product Image URL*
- Nutrition Data: *Attributes related to Vitamins, Calcium, Iron, Proteins, Calories / Energy , Carbohydrates, Sugars, Fat, Cholesterol, Sodium, Serving information*

The following data quality issues were discovered during the course of testing the system across the test products used in the pilot;

- Product Name and Product Description information was not consumer friendly in many instances.
- Product URL which is used to provide a link to the product's website, was mostly not available and when provided linked to the brand-owner's main website, rather than to a product-specific website.
- Product Image URL was mostly not available. When available, had varying image sizes and resolutions providing an inconsistent consumer experience when the images were displayed in the mobile application.
- Nutrition codes were not well defined in many cases leading to additional data harmonization and mapping across multiple product records. In some cases usage of nutrition codes was inconsistent across product records from the same brand. Some examples of different nutrition codes provided for the same nutritional attribute are:
  - Carbohydrates: CHO-, CHO, Carbohydrate, Carbohydrates
  - Sugars: SUGAR-, SUGAR, Saccharin, Sucralose, Sugars
  - Sodium: NA, NACL, Sodium
  - Fat: FATNLEA, FAT
  - Energy: Calories, ENER, ENER-
  - Vitamin A: VITA, VITA-
  - Vitamin C: VITC, VITC-

★**Recommendation**: Develop data quality requirements for B2C applications for brand-owners, data pools and data aggregators.

5.1.2. GDSN Datapools

The following data mapping issues were discovered during the creation of the B2C Data canonical schema, while defining the mapping of B2C data attributes to their equivalent GDSN master data attributes

- Due to the presence of multiple candidate GDSN master data attributes, there was ambiguity on which GDSN attribute mapped directly to the Product Name, Product Description and Serving Size information B2C Data attributes. As a result different GDSN Datapools provided data from different GDSN master data attributes for the same B2C attribute.
- Additionally, the related data fields in GDSN for Product Name and Product Description are populated with supply chain / B2B / Point of Sale data and these descriptions do not necessarily correlate to consumer information needs.
- GDSN master data standards uses the UN INFOODS code list as the standard for nutrition codes used for description of nutrition attributes. This is a very complicated and often misinterpreted code list leading to ambiguities. The data quality issues raised above is likely a result of this complexity.

★**Recommendation**: Develop implementation guidance for populating B2C product master data via GDSN.

5.1.3. Private Label

An RCN-8 is an 8-digit code beginning with GS1-8 Prefixes 0 or 2, as defined in the GS1 General Specifications. These are reserved for company internal numbering, and are not GTIN-8s. Many retail chains have been using RCN-8s for private label / own label products. Since these numbers are reserved for internal numbering, they cannot be used in the open supply chain for product identification. There is nothing stopping two different retailers from using the same RCN-8 to identify their unique private label products, as such RCN-8s are not globally unique. The EPC system and subsequently the ONS are based on the principle of globally unique identification and as such RCN-8s are excluded from use. The net impact is that private label products with RCN-8s cannot be used in a TSD framework without violating GS1 system principles.

During the course of the pilot, 4 products with RCN-8s were provided by a retail brand for their private label products. Since the B2C pilot was executed as a closed system, the RCN-8s were treated as globally unique GTIN-8s as it was guaranteed that there would be no conflicts. However, this could never work in a production TSD framework.

This is one of the big challenges for use of product bar codes in consumer-facing applications. The "closed system" is no longer closed, as products are not only scanned within the confined store environment which would otherwise serve to disambiguate. Retailers need to provide B2C information to their shoppers and IAPs for private label brands that may be using RCN-8s.

★**Recommendation**: Provide guidance on how private label products marked with numbers for internal use (such as RCN-8) can be used in the global TSD framework.

5.2. Technical Architecture

5.2.1. Object Naming Service (ONS)

The pilot utilized a technical setup of the ONS wherein the zone maintenance of all ONS records including the local ONS was provided by the root ONS operator. Ideally, the ONS system should be configured based on a nameserver delegation model wherein the peer ONS root nameserver under GS1's (Global Office and Member Organization's) authoritative control provides resolution at the GTIN company prefix level and the local ONS nameserver under the brand owner's authoritative control provides resolution at the GTIN item reference level. Although the current ONS setup works given the limited scope of the pilot, additional considerations need to be made for a commercial scale deployment of ONS in a future GS1 TSD framework. The current setup would lead to number of DNS issues related to zone maintenance, domain control and such. If the current 'centralized' ONS model is used moving forward, it must fallback to a setup wherein the zone maintenance of the local ONS is under the control of the brand itself, either directly or provided by their respective local GS1 MO or DNS provider as a hosted service

Additionally, based on concerns that are wholly outside the scope of the B2C pilot, a Federated ONS approach is preferred. It provides for distributed control in a system designed for the wider "Internet of Things", mitigating the issues concerning a single resolution root associated with the current 'centralized' ONS model

The GS1 Federated ONS initiative provides additional insight into the issues related to the current 'centralised' model of ONS and should be a consideration as a solution for future adoption.

★**Recommendation**: Implement Federated Object Naming Service (F-ONS) as the registry function allowing a product's Global Trade Item Number (GTIN) to be matched with the location aggregated product information.

5.2.2. Gateway Service

One of the key learnings from the B2C pilot was that the gateway service as a standalone entity is not needed in a future TSD framework for a number of reasons:

- It is a single point of failure in the current architecture of the TSD framework and would never be realised in a commercial deployment for practical and technical reasons.
- B2C product information needs are almost always local to the IAPs given target market of coverage. IAPs information needs could be serviced more effectively by querying their local data aggregator directly.
- The local data aggregator can effectively perform the functions of a gateway service and directly service IAP requests, eliminating the need for an additional step in the GTIN resolution process when product information is locally available most of the time.
- This new approach also results in performance benefits from the overall system by limiting the need for full ONS resolution to only when product information is not available at the local data aggregator.

★**Recommendation**: Ensure local data aggregators can directly service IAP requests, so eliminating the need for a gateway service managed by GS1. This approach is shown in the architecture vision in section 6.

5.2.3. GTIN-EPC Interoperability

In the pilot, the Gateway Service played the role of an ONS client. It performed the critical function of ONS lookups to resolve a GTIN to its corresponding data aggregator web service URL. This requires deriving the correct ONS domain name corresponding to the GTIN. However due to a parsing issue related to EPC – bar code interoperability, there is no direct way of deriving this reference from the GTIN. To get the correct ONS domain name reference for resolution, an ONS client has to make multiple ONS queries with varying lengths of company prefix for the GTIN in consideration, until ONS returns a conformed successful or failed response. The need to make

multiple ONS queries introduces additional latency in the network used by the TSD framework and results in increased response times. This can be minimized if not eliminated, if a proper software solution was in place to address the GTIN/GCP parsing issue.

★**Recommendation**: Develop and deploy a standard GS1 solution to ensure bar code numbers can be easily and unambiguously resolved by ONS.

5.2.4. Security

Pilot participants expressed concern over the lack of guidelines for security of data in the network. Given the limited scope of the pilot, security considerations were considered out of scope for the pilot. Security concerns arising from lack of authentication, authorization and encryption were masked by virtue of the pilot being a closed system. Any commercial product / production system should have well defined security controls to mitigate the risk arising from the lack of it.

★**Recommendation**: Develop guidelines addressing all aspects of security in the TSD framework.

5.2.5. Performance

The B2C data model utilised in the pilot was well defined and met the business requirements outlined for the pilot. However, during the course of systems development and testing of the TSD framework two attributes were identified as important considerations for future standardization.

- Record Last Updated: The date when the B2C data was lasted updated by the Brand Owner
- Record Time To Live: The validity period of the data record

It was discussed that the combination of the two fields above can be used to improve overall system performance by enabling support for caching of data in the network used by the TSD framework.

★**Recommendation**: Consider options to increase system performance by supporting caching of data in the TSD framework.

5.3. Data Usage

5.3.1. IAP

The pilot utilised a simplified XML data format consisting of key-value pairs that was provided by the Gateway Service to the mobile application. Future versions of the infrastructure should consider JavaScript Object Notation (JSON) as an additional mark-up format, which is often used in mobile development and which simplifies data integration into mobile applications.

★**Recommendation**: Use an additional data format (such as JSON) that simplifies integration into mobile applications by IAPs.

5.3.2. Target Market

Consumer needs for product information are highly dependent on:

1. **Target market where product is sold**. Internet Application Providers will almost always provide information in the language of the target market. In addition a number target markets have multilingual requirements. Nutritional measurements are also expressed in the local measurement system and in compliance with local regulations for declaration of nutritional information.
2. **Language of the consumer**. A consumer may be travelling outside their country of origin but expect to receive product information in their own language.

One of the shortcomings identified in the pilot was that although the product GTINs were successfully resolved and authentic product data was available, the product information did not provide a target market indicator making localization of product information very difficult.

The pilot identified a need for brand owners to provide target market specific language translations and measurement units to avoid any ambiguity during the localization process.

★**Recommendation**: Better understand B2C requirements related to target market and provide appropriate standards and guidelines.

6. Next steps

The pilot can be deemed a success and the learnings from it are an opportunity to build the future foundations of a GS1 B2C Trusted Source of Data framework.

GS1 and relevant stakeholders will now focus on the deployment of the Global B2C TSD Framework by development of global standards and guidelines and a common operational model. Local aggregation in some cases has already begun and will conform to the global framework as important aspects are deployed. The goal will be to work towards the architecture vision shown below.

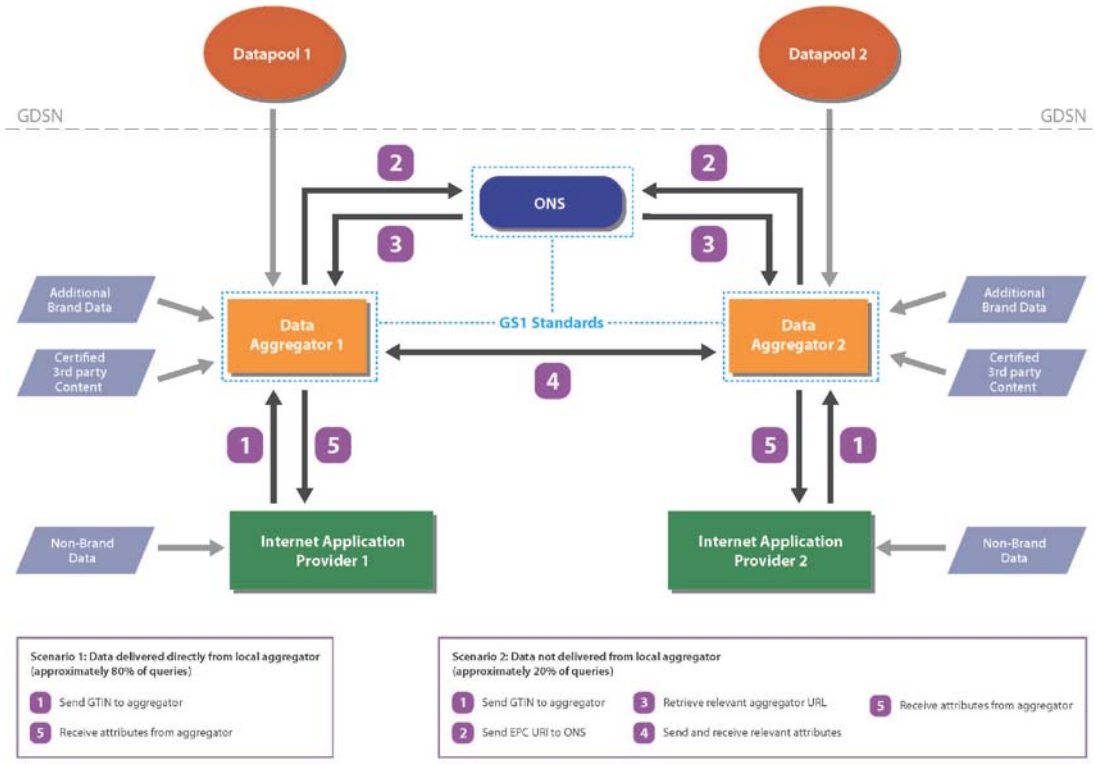

Figure 3: Trusted Source of Data Framework vision

For more information about the pilot and the current status of the GS1 B2C Trusted Source of Data project, contact Cameron Green at cameron.green@gs1.org.