THE INTERNATIONAL SOCIETY OF PRECISION AGRICULTURE PRESENTS THE
13th INTERNATIONAL CONFERENCE ON
PRECISION AGRICULTURE
July 31-August 4, 2016 • St. Louis, Missouri USA

Toward geopolitical-context-enabled interoperability in precision agriculture: AgGateway's SPADE, PAIL, WAVE, CART and ADAPT

Applegate, Doug B.[10]; Berger, Aaron W.[4]; Berne, Daniel T.[9]; Bullock, Rob[14]; Craker, Ben E.[2]; Daggett, Dennis G.[12]; Ferreyra, R. Andres[1,*]; Gowler, Andrea[14]; Haringx, Shannon C.[16]; Hillyer, Charles[17]; Howatt, Todd[2]; Nef, Bart K.[5]; Nieman, Scott T.[8]; Reddy, Linga T.[7]; Rhea, Stuart T.[1]; Russo, Joseph M.[19]; Sanders, Patrick[15]; Schultz, Eric D.[13]; Shearouse, Timothy W.[7]; Stelford, Mark W.[11]; Tevis, Joe W.[18]; Wilson, Jeremy W.[6]; Wilson, James A.[3]

[1]Ag Connections LLC, Murray, KY, United States; [2]AGCO Corporation, Duluth, GA, United States; [3]AgGateway, Washington, DC, United States; [4]AgSense LLC, Huron, SD, United States; [5]Campbell Scientific Inc, Logan, UT, United States; [6]CropIMS, Effingham, IL, United States; [7]John Deere ISG, Urbandale, IA, United States; [8]LAND O'LAKES, INC., Arden Hills, MN, United States; [9]Next Chapter Marketing, Portland, OR, United States; [10]Praxidyn, Oakland, IA, United States; [11]Premier Crop Systems, LLC, Des Moines, IA, United States; [12]ProAg Management, LLC, Amarillo, TX, United States; [13]Schultz Consulting, Des Moines, IA, United States; [14]Software Solutions Integrated, LLC, Shelbyville, IL, United States; [15]SST Software, Stillwater, OK, United States; [16]Syngenta, Greensboro, NC, United States; [17]Texas A&M AgriLife, Amarillo, TX, United States; [18]TOPCON, Minneapolis, MN, United States; [19]ZedX, Inc., Bellefonte, PA, United States
* Corresponding author

A paper from the Proceedings of the
13th International Conference on Precision Agriculture
July 31 – August 4, 2016
St. Louis, Missouri, USA

Abstract. *AgGateway is a nonprofit consortium of 240+ businesses working to promote, enable and expand eAgriculture. It provides a non-competitive collaborative environment, transparent funding and governance models, and anti-trust and intellectual property policies that guide and protect members' contributions and implementations. AgGateway primarily focuses on implementing existing standards and collaborating with other organizations to extend them when necessary.*

In 2010 AgGateway identified interoperability in field operations (e.g., planting, harvest) as a major challenge: the myriad proprietary, incompatible data formats among different machine and implement control systems (MICS) and farm management information systems (FMIS) leads to user frustration and operational inefficiencies.

The SPADE project was AgGateway's first field operations effort. SPADE has generally followed an Agile Methodology for each field operation: collect expert-mediated user stories; formalize them into business process models and use cases; identify data requirements; perform gap analyses of existing standards, and propose extensions thereto. A key SPADE outcome was the identification of a set of flexibly-defined common "documents" that support farm business processes: Crop Plan, Observations and Measurements (O&M), Recommendation, Work Order, and Work Record.

SPADE identified Reference Data: datasets to distribute across the whole industry so different stakeholders can interpret shared documents the same way. This includes names and identifiers of seed varieties, crop protection products, active ingredients, etc. SPADE implemented proof-of-concept (PoC) application programming interfaces (APIs) for machinery and product Reference Data, and also implemented a PoC index API providing one-stop access to a distributed system of standardized Reference Data sources hosted by manufacturers and third-party data providers. Other SPADE contributions include defining concepts such as "OK to Spray", contributing terms to AgGlossary.org, and working with the Agricultural Electronics Manufacturers' Foundation (AEF) to expand the widely-respected ISO 11783-10 MICS-FMIS communication standard's ability to associate globally-unique identifiers to its own locally-scoped identifiers.

CART is a SPADE sub-project. Its focus is to support grain movement and testing processes, aligning with (and expanding) the AgXML Consortium's CommodityMovement and QualityCertificate messages.

WAVE is another SPADE sub-project. Its scope includes mobile and fixed asset management. It seeks alignment with the new ISO 15143-3 mobile asset telemetry standard. Its fixed asset scope (e.g., grain bin temperatures) also seeks existing-standards alignment, and dovetails with O&M work in PAIL and various SPADE teams.

PAIL is SPADE's irrigation-specific sibling. Given the lack of existing irrigation data exchange standards, PAIL is proposing one (via ASABE) that documents irrigation operations and related O&M, including web service-mediated data-exchange schemas.

SPADE successfully explored the feasibility of an open-source format-conversion toolkit, leading to what is now ADAPT. ADAPT's scope includes a common object model derived from SPADE / PAIL requirements including ISO 11783-10 compatibility, a format-conversion framework, and an architecture where manufacturer-specific "plugins" convert between proprietary formats and the common model. ADAPT supports FMIS-to-FMIS communication, including the nuanced business processes not covered in ISO 11783-10.

Internationalization is an important aspect of this work; several conflicting requirements (universality vs. geopolitical-context-specificity, and controlled vocabulary vs. extensibility) were reconciled creating a ContextItem class and drawing from AEF's experience managing controlled vocabularies separately from the ISO 11783-10 data model.

Keywords. *Farm management information systems, FMIS, precision agriculture, field operations, irrigation, grain, telematics, workflow, farm data, standards, geopolitical context.*

Introduction

Field operations such as planting, irrigation and harvest are grower business processes. These processes vary in detail from grower to grower, but there is value in defining a broadly-applicable core set. To that effect, and in order to support the implementation of quality management systems for production agriculture, ISO Standard 22006 (ISO, 2009) provides a list of crop-production processes and sub-processes.

Crop-production processes must be supported with data exchange, both for coordination among the different actors directly involved in each process, and for communication with information consumers such as trade partners and regulatory agencies.

A major problem in agriculture, however, is a lack of standards to govern this business-process-centric data exchange. A well-respected standard, ISO 11783-10 (ISO, 2015), does exist for a subset of these processes; specifically, for communication between machine and implement control systems (MICS) in the field and farm management information systems (FMIS).

Whereas ISO 11783-10 provides a comprehensive format for representing data being exchanged between an FMIS and machines such as tractors, sprayers and combines, it has limitations:

- It is limited to MICS-FMIS communication in the context of field operation execution, and does not support FMIS-FMIS transactions that may involve additional "documents" (e.g., a recommendation from an agronomist, or data needed to complete a regulatory report) that are not machine-specific and therefore not covered by the ISO 11783-10 format.
- It initially supported locally-scoped identifiers only.
- It is not well-suited for some field operations such as irrigation with center pivots, where irrigated areas have complex geometries.
- Its adoption is not yet universal, and there exists a multitude of proprietary formats for machinery and FMIS that are inherently incompatible with it.

The lack of interoperability in agricultural field operations is not just a problem of a lack of common data formats or syntax. There has also been a lack of a shared understanding, or semantics, among the different industry actors involved in field operations. This can take the form of using multiple terms to refer to the same concept, or using the same term to refer to multiple concepts. Some anecdotal examples may help illustrate this conflict in semantics:

- Agreeing on the meaning of the word "field" is often difficult among different industry actors;
- The names "task", "job", and "work order" have often been used interchangeably, but mean subtly different things for different manufacturers.

There is also a lack of shared *reference data*: data that unambiguously identifies products such as seed, crop protection, and equipment brands and models. When two actors exchange data, say, a report of product usage, it's important that both participants interpret that the products involved were the ones actually used. Given the complexity of some product names, particularly in crop protection, this common interpretation cannot be easily attained without some form of *unique identifier* such as a universally-unique identifier or UUID (Leach, 2005) shared by participants in the data exchange process. This is especially important because in many jurisdictions in the world seed and crop protection products must be used in strict accordance with requirements specified in their label: genetically-modified varieties often require refuges to prevent the development of pest resistance, chemicals may have strict application protocols including personal protective equipment, maximum allowed active ingredient loads, and so forth. In a context where application data for these products are being exchanged, it is critical that products be identified correctly, and unambiguously. Current business practices do not support this requirement. Specifically, many forms of equipment do not support unique product identifiers, and several companies provide product label data in the marketplace with each using their own proprietary identifiers, formats, and delivery mechanism or application programming interface (API).

AgGateway (www.aggateway.org) is a nonprofit consortium of 240+ businesses working to promote, enable and expand eAgriculture. It provides a non-competitive collaborative environment, transparent funding and governance models, and anti-trust and intellectual property policies that guide and protect members' contributions and implementations. AgGateway primarily focuses on implementing existing standards and collaborating with other organizations to extend them when necessary. For the past several years, groups within AgGateway have been addressing interoperability problems in agricultural field operations, in addition to a decade-long effort to do the same in agricultural supply-chain management. The goal of this paper is to present AgGateway's field operations work and its significance to a wider audience. Specific objectives are to:

- Frame the field operations interoperability problem,
- Introduce the AgGateway field operations – themed projects (SPADE and PAIL), the process behind them, additional telematics (WAVE) and grain handling (CART) efforts, and their deliverables (process models, Core Documents, Reference Data APIs, etc.)
- Introduce related infrastructure development, especially the ADAPT open-source framework, and
- Present the problem of incorporating geopolitical-context-dependent data into a generic field operations data model, and the SPADE/PAIL/ADAPT proposed solution: the ContextItem system.

First Steps and a Vision

There have been efforts in the direction of providing common formats for interoperability in agriculture. For example, the FODM (Field Operations Data Model), a proprietary data model and set of format-conversion tools (Macy, 2002), is still widely used today. It is, however, proprietary, not specifically ISO-compatible (pre-dating much of the recent ISO development), and does not include many business process documents used in production agriculture.

An industry-supported Reference Data source for seed and crop protection products exists as well: the Ag Industry Identification System or AGIIS (AgGateway, 2015). AGIIS contains over 140000 product identifiers in the form of Global Trade Item Numbers or GTINs (GS1, n.d.) and is a powerful resource for supply-chain operations. However, it lacks some data important to field operations, such as package-agnostic product-specific unique identifiers, active ingredient loads, and so forth.

In this context, AgGateway created its Precision Ag Council in 2010. This group, which has grown to include over 100 companies in the agriculture industry, started working along four distinct lines:

- An agricultural glossary (www.agglossary.org), meant to bring together agricultural terminology from different sources as an educational resource and discussion-support tool.
- An interest in Reference Data and the problem of developing a distributed system for sourcing product identifiers.
- A Field Operations Working Group, which brought together the companies that were interested in field operations interoperability, and
- A Telematics and Asset Management Working Group, for those parties interested in the non-task-specific management of fixed and mobile assets.

The Field Operations Working Group eventually chartered a series of projects, SPADE1, and SPADE2, which dealt with various field operations and with Reference Data; the SPADE3 project subsequently included a telematics component, called WAVE. An interest in irrigation resulted in a Water Management Working Group within the council, and the PAIL1 and PAIL2 projects followed. Likewise, a desire to integrate grain operations motivated, through interactions with the Grain and Feed Council, another SPADE3 component: CART.

Figure 1 portrays the relationships among the different AgGateway councils, working groups and projects discussed in this paper, as well as the data exchange standards they involve. Note how the work includes interacting with multiple standards bodies: the American Society of Agricultural and Biological Engineers (ASABE), the Association of Equipment Management Professionals (AEMP), the International Organization for Standardization (ISO), and so forth. This is consistent with

AgGateway's usual approach of seeking to implement and enhance existing standards, only creating new ones where something does not already exist for the desired purpose.

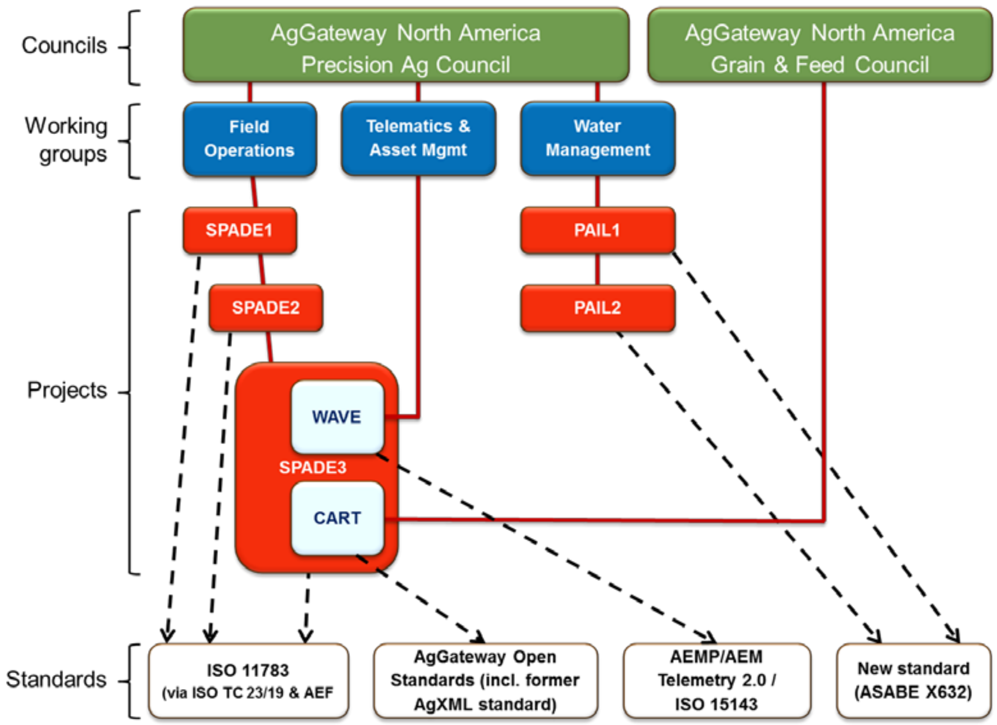

Figure 1: Relationships among AgGateway North America (a rebranding of the original AgGateway organization following the creation of AgGateway Global Network) industry segment councils, working groups, projects and standards. Councils may be organized into working groups (three of the Precision Ag Council's groups are shown). These working groups can propose and organize projects for the council to charter. The projects in turn can influence, or contribute to, standards.

In much the same way that AgGateway's supply-chain work has streamlined the exchange of information among manufacturers, distributors and retailers of various agricultural inputs, the AgGateway Precision Ag Council's vision (Figure 2) addresses the problem from the other direction: starting from the grower and the various sources of data with which the grower works (modeled in SPADE, PAIL and WAVE). The aforementioned CART work, detailed further below, represents one of the spaces where the top-down and bottom-up approaches meet.

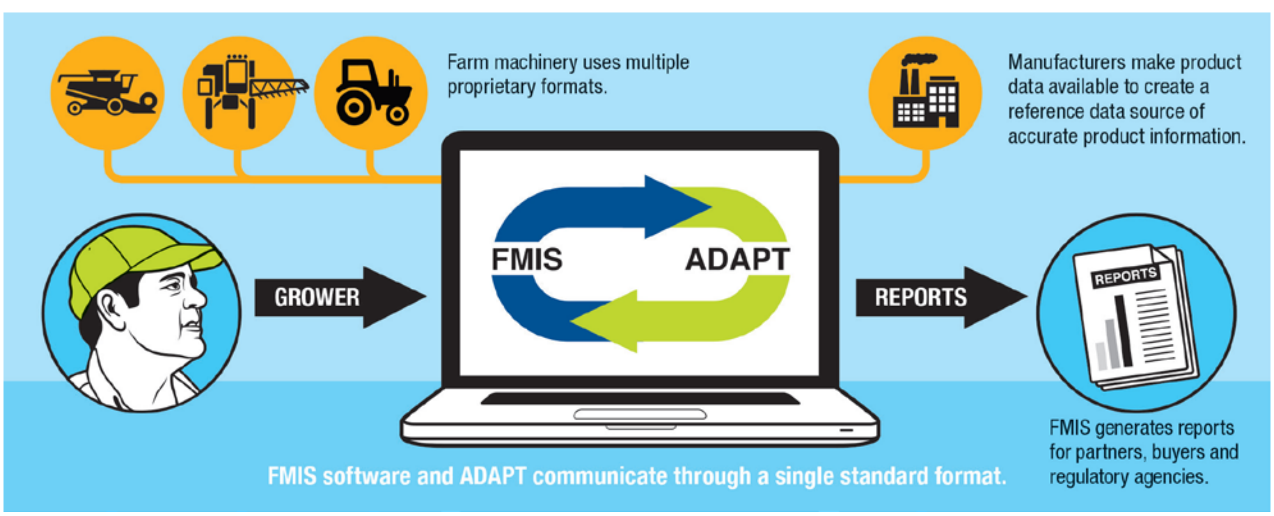

Figure 2: AgGateway's vision of how to tackle interoperability problems in field operations through the use of a common object (data) model, data format-conversion tools, and Reference Data.

The Development Process

The SPADE and PAIL projects have generally followed an Agile Methodology (Agile Alliance, 2013), prioritizing individuals and interactions, emphasizing the need for usable deliverables, promoting collaboration, and responding to change.

The projects created separate *product teams* for each field operation (e.g., seeding, mechanical harvest, etc.), and each corresponding product team loosely followed the development process shown in Figure 3:

- Collected user stories (Agile Alliance, 2016) from growers and other domain experts (Table 1).
- Modeled the processes described by the stories using BPMN, the Business Process Modeling Notation (Object Management Group, 2011).
- Where the groups deemed it necessary, they translated the process models into use cases (Jacobson et al., 1992) for the benefit of software companies.
- Data requirements were identified through joint work by the domain and technical experts.
- The technical experts then looked for, and proposed solutions to, gaps in existing standards (ISO 11783, primarily).
- Publication followed as a final step.

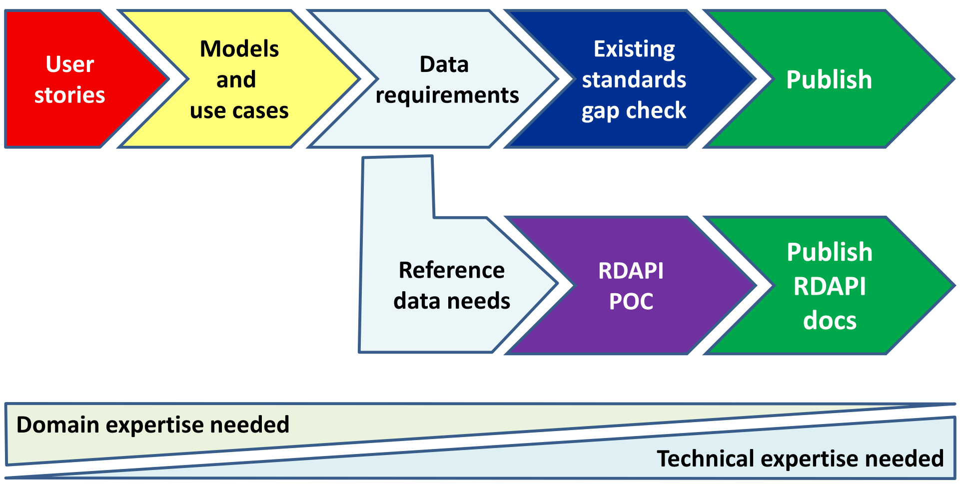

Figure 3: The development process followed by the SPADE and PAIL product teams.

The effort required collaboration between domain experts (e.g., agronomists, sales representatives) and technical experts (e.g., information architects, developers), with the degree of participation of each kind of expert changing according to the task at hand (Figure 3, bottom): domain expert input was critical in the early stages of the process, whereas technical experts were primarily engaged later.

Reference Data was a special case of data requirements work. Reference Data requirements were identified (and iteratively refined) throughout the projects. This was followed by the design of RESTful application programming interfaces or APIs (Fielding, 2000) to deliver the Reference Data to users (also a subject for iterative refinement throughout), and the documentation thereof. Proof-of-concept implementations currently exist for Reference Data APIs.

Table 1. Example user stories from the PAIL project.

As a/an	I want to…	So that I can
Grower	create an irrigation Work Order	be sure the Irrigator knows how much water to apply and where to apply it.
Irrigator	use the irrigation Work Order to send a command to the irrigation system controller	begin and end the irrigation as planned, or modify as field conditions change.
Data Provider	store a Work Record of what actually happened during the irrigation event	provide a record as requested from an authorized user.
Consultant	retrieve a Work Record of the irrigation event	use the data as input for the next irrigation Recommendation.
Grower	store and retrieve a Work Record	use it as input for planning next season's crops and field operations, and provide reports, as necessary, to regulators and/or insurance providers.

Business Processes and the Documents that Support Them

As mentioned earlier, business process modeling has been an important aspect of the SPADE/PAIL effort. An example of the translation of user stories (Table 1) to a process model (Figure 4) follows. A detailed description of BPMN, the notation used for such modeling, is out of the scope of this paper. However, a quick introduction follows, sufficient to understand the example:

- Different *actors* (Grower, etc.) are represented by the rectangular horizontal *pools*.
- The processes carried out by each actor are contained in the corresponding actor's pool, and are represented as a connected sequence of *activities* (rounded rectangles).
- Processes begin, end, and sometimes are paused by *events*, shown as circles in the diagrams.
- There are different kinds of events, such as actions triggered by time (shown with a clock-face icon) or by receiving a message (shown with an envelope icon).
- Communication among pools happens through *messages*. Some of these messages (Work Order, Recommendation, Work Record) belong to a group of "Core Documents" that appear recurrently in SPADE and PAIL. The Core Documents will be discussed in greater detail later.
- The flow of a process can fork, depending on the outcome of an activity. The places where flows diverge (or converge) are shown with *gateways* (diamond shapes, not shown in our example.)

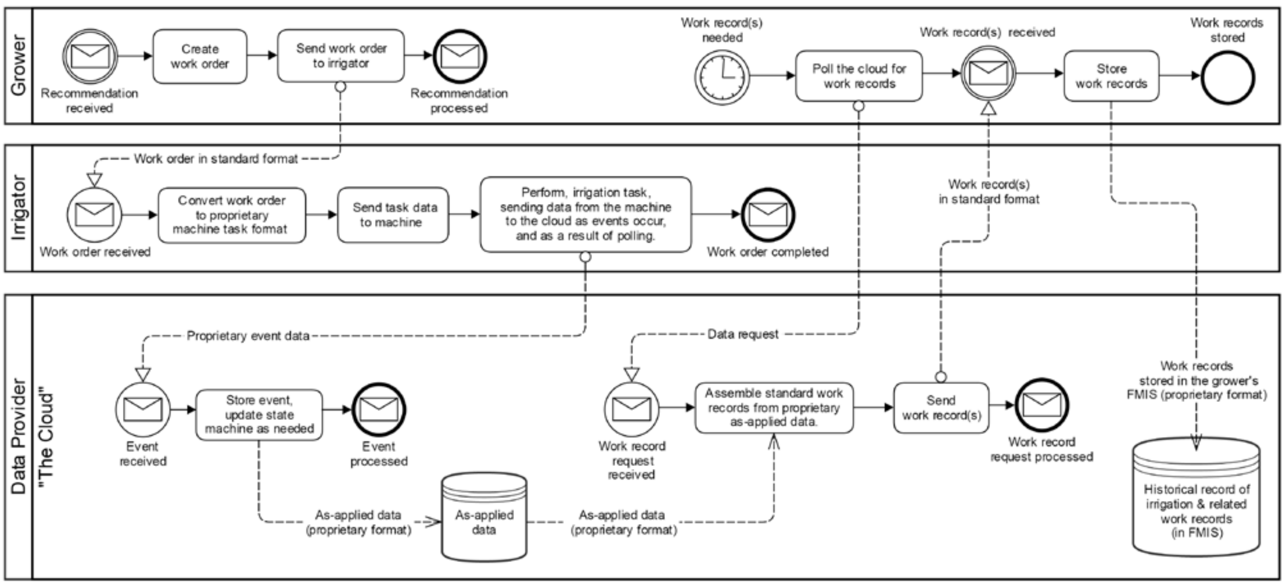

Figure 4: Business Process Modeling Notation (BPMN) diagram of example, taken from the PAIL project.

The example above shows five different processes involved in irrigation operations:

- Grower creating a Work Order (from a received Recommendation) and sending it to the Irrigator.
- Irrigator executing a Work Order received from the Grower.
- Data Provider storing event data received during the irrigation field operation.
- Data Provider assembling Work Records from stored event data, and sending them to the Grower upon request.
- Grower requesting Work Records from a Data Provider and storing them in an FMIS.

Note how the BPMN diagram of Figure 4 converts the simple requirements laid out in the user stories of Table 1 into more formal, sequential descriptions of *what* is done, *how* it happens, *who* (actor) is involved, and how information flows between them.

Although Figure 3 implies that the steps in our approach were sequential, there was plenty of iterative revision, opportunistic "borrowing" of structures, and a constant search for commonalities across different SPADE/PAIL product teams. The following section will describe the emergent model for representing the data needed to support field operations business processes: The Core Documents.

AgGateway's Core Documents for Field Operations

Growers and their partners use multiple *documents* to exchange field operations information as part of their business processes. There has been work done on standardizing farm processes (ISO 22006), but the documents and many of the terms used within them have yet to be unambiguously defined. The Core Documents (enumerated in Table 2) were an outcome of the SPADE and PAIL projects: they define the data exchanged during field operations (i.e., Figure 3, third stage). The definitions are quite flexible, in view of the myriad ways in which different growers implement their record-keeping in response to regionally-specific regulatory requirements, particular characteristics of their markets or farming operations, and personal preference.

Table2: The Core Documents and Supporting Information

Document Name	Type	What It Conveys	Actor Involved
Crop Plan	Strategic	A high-level document describing how a crop will be grown on a given piece of land during a crop season: "This is how we're going to grow this crop this season."	Grower, or other actor involved in the strategic planning for the field operations.
Observations and Measurements	Tactical/ Predictive	A document containing data measured / observed in the field: "This is what's happening (or what we think might happen) in the field."	Crop scout, remote observation or a person tasked with monitoring conditions in the field.
Recommendation	Tactical	"This is what I recommend we should do" This document is not always acted upon; it is acted upon via a work order, upon approval.	An individual, such as a consultant or agronomist, with the expertise / licensing necessary to recommend a course of action.
Work Order	Tactical	"This is what we are going to do."	An individual with authority to order the work done.
Work Record	Tactical/ Historical	"This is what we actually did in the field."	May be automatically generated; otherwise, an operator that performed the task.
Reference Data and Setup File (Not core documents)	All	"This is the common information we need in order to set up and support accurate and efficient data exchange."	Grower, or other actor involved in managing the grower's production data.

The documents rarely stand on their own: they exist as part of a complex network of decision-making processes within a grower's operation. Core Documents are linked through various relationships: causal (e.g., a recommendation *informs* a work order), contextual (e.g., a product label helps identity an insecticide used in a Work Order), and compositional (e.g., photographs and sound files can form part of a scouting report / observations & measurements document). Figure 5 shows the Core Documents, the causal relationships among them, and the contextual relationships afforded by Reference, Setup and Configuration data.

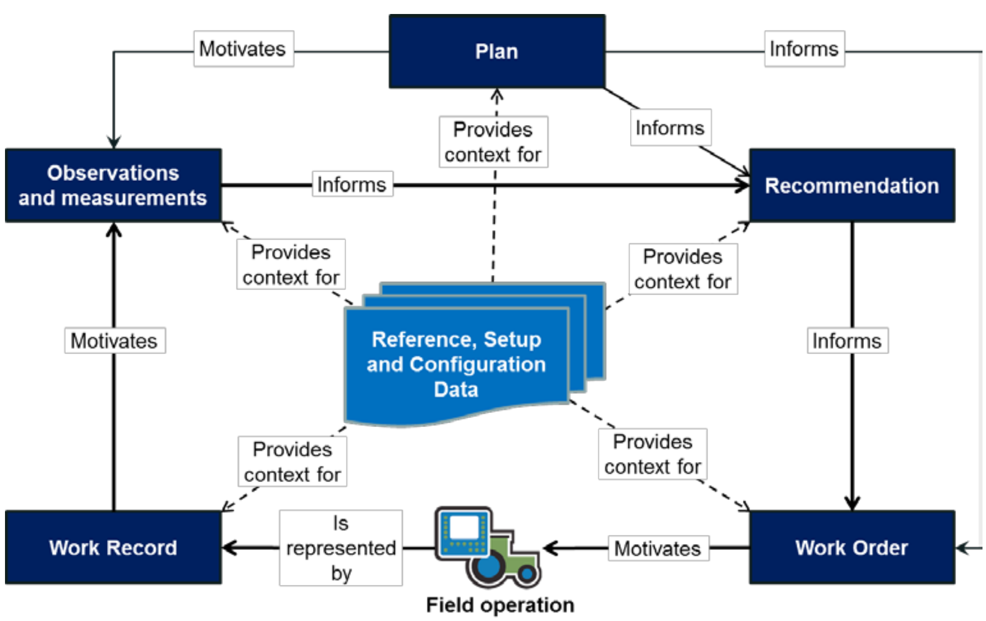

Figure 5: AgGateway's Core Documents for Field Operations, and the relationships among them.

Examples of the Core Document relationships shown in Figure 5:

- The Crop Plan informs or motivates the other documents.
    - o Example: a grower creates an annual plan outlining crop/variety allocation, crop fertility, irrigation, and crop protection strategies for a set of fields. This crop plan will be consulted when planning subsequent operations ("Did I budget for an expensive fungicide?")
- Observations and Measurements (O&M) inform Recommendations.
    - o Example: a high insect count observed by a scout indicates the need to spray.
- Recommendations inform Work Orders.
    - o Example: an agronomist recommends irrigating because the corn will flower soon.
- Work Orders motivate field operations.
    - o Example: A grower purchases a fungicide from a retailer and requests its application.
    - o Example: A grower communicates to an operator (irrigator actor) that a field must be irrigated with a certain depth of water over a certain period of time. (See Figure 4)
- Field operations are represented by Work Records
    - o Example: A combine harvester's controller creates a record of grain yield across the field.
- Work Records motivate Observations and Measurements
    - o Example: A crop scout goes out to the field to determine whether there are still symptoms of water stress in a crop following an irrigation operation.

The previous section described the Core Documents and the relationships among them. Figure 6 shows an example of their exchange as part of the Grower's business processes, as follows:

- The Grower shares the Crop Plan with a Crop Advisor and a Crop Scout (O&M) service.
- The Grower also shares a historical record of Work Records and O&M with the Crop Advisor.
- The Crop Scout makes a Recommendation informed by the Crop Plan, the historical record, and fresh O&M.
- The Grower, informed by the Recommendation, orders a course of action through a Work Order that is sent to a machine/operator that executes the field operation. (See Figure 4.)
- The Work Record is returned to the Grower; e.g., through a web service associated with an equipment manufacturer or other data provider. (See Figure 4.)
- The Grower processes the Work Record, creating a report shared with a regulator or value partner (e.g. a banker).

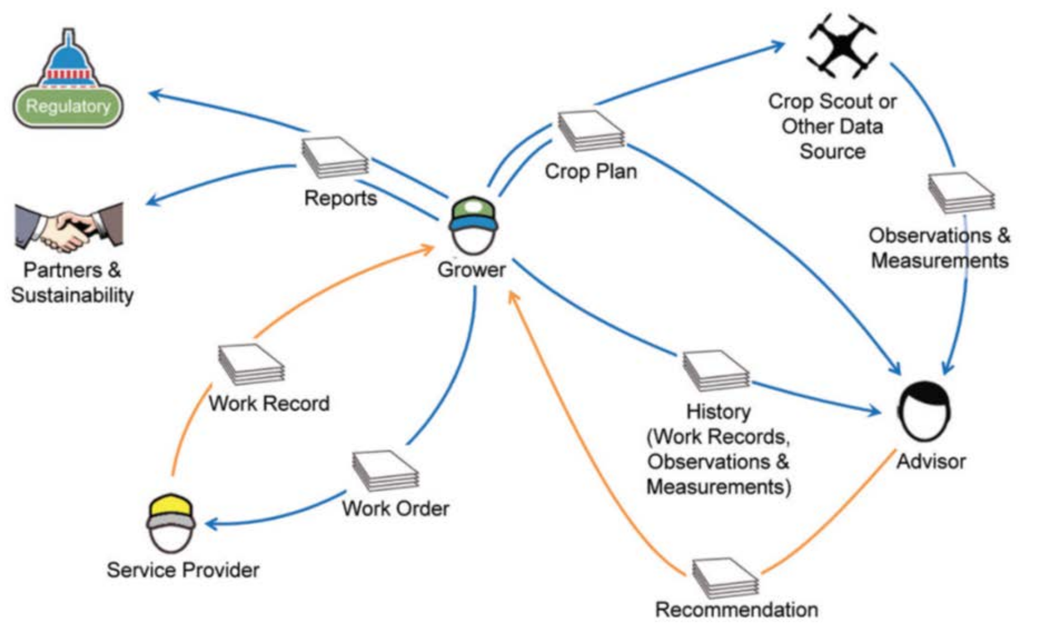

Figure 6: Examples of document flow.

PAIL: Standards for Irrigation Data Exchange

The purpose of PAIL is to provide an industry-wide format to enable the exchange and use of data for precision irrigation management systems. The data are currently stored in a variety of proprietary formats; the objective is to develop a common language that can enable integration of these disparate sources of water management information.

As noted above, AgGateway's Core Documents are designed to support field operations processes. In the case of precision irrigation, Observations and Measurements are captured and sent to an FMIS, or a Decision-Support System (DSS). The data inputs may include soil moisture, atmospheric, plant, and other field measurements that impact the decision on when and how much to irrigate. This data is provided by equipment such as soil moisture sensors, field weather stations, or crop-related sensors. Derived weather data from commercial or governmental weather providers, such as CMIS and AgriMet, add to the knowledge of site-specific weather data and forecasts. PAIL's data format is based on, and extends, the ISO 19156 standard for Observations and Measurements (ISO, 2011).

Another part of the PAIL data standards involves all of the activities associated with the irrigation operation itself, i.e., the application of water through the irrigation system. This includes, but is not restricted to, management-level communications and record-keeping. Once the Observations and Measurements are taken into account, together with the Crop Plan, an Irrigation Consultant can make a Recommendation for the grower. When approved or modified, the Recommendation is documented as a Work Order and sent to the irrigation system's controller. The conversion from the PAIL-formatted Work Order to a hardware-specific language can be accomplished either by a manufacturer individually or through a PAIL plug-in designed in the context of the ADAPT framework. The actual execution of the Work Order is left up to the proprietary processes of the individual manufacturers. However, once the irrigation event is completed, or even if it is interrupted, the irrigation consultant can send a request to the controller system and receive a record of how much water was applied over the designated area and when it happened. The Work Record can be used both for required reporting purposes and as a valuable historical record for the Grower and Consultant to use for planning future irrigation events.

SPADE, PAIL and Existing Standards

AgGateway is committed to leveraging, whenever possible, applicable existing standards such as ISO 11783-10, ISO 19156, GS1's GTIN and GLN, and OAGi's Chem eStandards. The SPADE and PAIL projects were consistent with this commitment.

For (non-irrigation-related) field operations data requirements, SPADE project participants instinctively looked to ISO 11783-10 (ISO, 2015). While it initially proved to be a sturdy point of reference, it gave way to the concept of Core Documents for field operations, which the SPADE project participants developed in a more top-down and comprehensive fashion with a commitment to align with and enrich ISO 11783-10 where applicable.

The PAIL project did not have ISO 11783-10, or anything like it, as an initial point of reference. This led to more early white-board activity than in SPADE. As the Core Documents developed in SPADE, the PAIL project participants took notice and began mapping their requirements against them as defined in the then-nascent ADAPT data model. Concurrently, the PAIL project participants intuited that their field-data-collection-specific requirements abstracted to a pattern they were certain must have been formally addressed by some standards group. This proved to be the case: they discovered ISO 19156 and engaged the lead editor thereof to explore the potential applicability to PAIL, and opportunistically, scouting in SPADE.

Given the success of applying AgGateway Open Standards (a rebranding of OAGi's Chem eStandards) to supply-chain operations processes, the SPADE and PAIL project teams naturally investigated their applicability to the field operations domain, concluding that they were not applicable to their work beyond a few data structures and concepts.

The end-result of all the inquiries into other standards left the Core Documents as the central data concepts of SPADE/PAIL, expressed precisely in the ADAPT data model, informed by ISO 19156, and compatible with ISO 11783-10 where applicable.

Grain Handling and CART

Grain-sector processes revolve around grain contracts, grain shipments, and settlements. Contracts specify quantity, quality, timing and price, along with premiums and discounts associated with exceeding or falling short of contracted baselines. The process of determining how a delivery compares to contract terms requires quality measurements and grain weight—or weights and grades. Historically manual and paper-centric, the industry decided that adopting electronic versions of contracts, weight certificates, quality certificates and settlements could improve the quality and speed of data-management processes while eliminating costly and error-prone manual work.

AgXML was formed in late 2000 by several large grain companies, joined by a number of software and service providers. AgXML's mission was to identify, develop, and implement standards and guidelines intended to bring electronic business efficiencies to grain and oilseed companies. Each company committed contributors to a project to achieve the mission, which was realized by documenting processes, identifying where information was exchanged in the context of each process, and specifying the structures and data types of the information contained in the abovementioned documents.

AgXML developed standards to define the electronic forms of those documents in an industry forum such that interoperability would be achieved. Furthermore, the industry defined another message, CommodityMovement, to tie together all the documents related to a grain shipment. AgXML subsequently released standards for bills of lading, commodity movement, contracts, contract pricing, quality certificates, weight certificates, rail rates electronic exchange, biofuels support, and settlements. AgXML also worked with the USDA's Grain Inspection, Packers & Stockyards Administration (GIPSA) to update AgXML's quality certificate standard to meet GIPSA requirements.

AgXML standards adoption is widespread among grain companies in the industry, enabling the benefits described above. GIPSA has endorsed the use of AgXML's QualityCertificate standard, and a number of software and service providers have implemented AgXML support in their solutions.

In 2011, AgXML member companies, observing AgGateway's success in driving industry implementation projects, reached out to establish collaboration with AgGateway. AgGateway established a Grain Council and AgXML's Management Committee agreed to conduct future standards-development activities within it.

Some time later, SPADE project participants identified a grower need for more effective electronic data exchange at the interfaces between harvesting equipment and transport vehicles through delivery to a grain elevator or storage bin. Concurrent with this discovery was a desire within AgGateway's Grain & Feed Council to identify an electronic connectivity project to move the grain industry forward. This led to precision-agriculture experts discussing their requirements with grain-industry experts. The CART product within the SPADE3 project grew out of those conversations.

Telematics and Mobile / Fixed Asset Tracking: WAVE

The media is extensively covering Internet of Things (IoT). This coverage typically predicts that enhanced connectivity and ubiquitous sensors will lead to fundamental shifts in almost every industry and potentially most aspects of everyday life. The agricultural industry is beginning to embrace this transformation, and in many respects is actually leading the charge for more connected and intelligent machines. A recent review on the topic (Porter and Heppelmann, 2014) showed a progression from simple products to a system of systems (Figure 7).

Many tractors and other agricultural machines were originally strictly mechanical products (Figure 7, 1st stage) that could only operate at peak efficiency with a very skilled driver. As sophisticated electronics became common and reliable, equipment manufacturers integrated it into controllers, creating "smart products" (Figure 7, 2nd stage), machines that performed better and with less operator input, with more fuel-efficient engines and electronically-controlled hydraulics and transmissions that enabled quicker response and improved operational efficiencies.

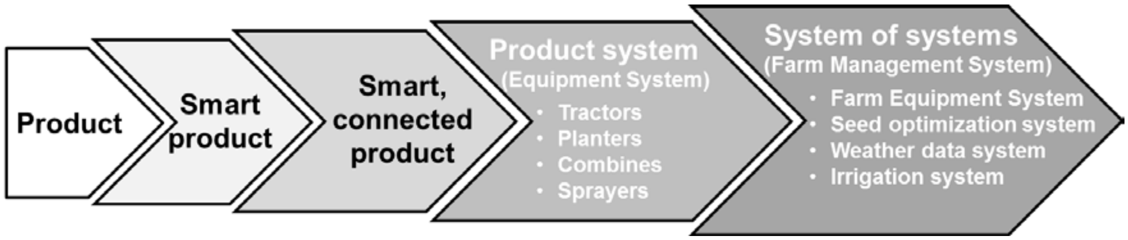

Figure 7: Porter and Heppelmann's (2014) model of the evolution of connected products.

The adoption of GPS technology and cellular communications unlocked the ability to remotely monitor the position, status, health and performance of "smart, connected" machines (Figure 7, 3rd stage). Today many growers use this information to improve their operations through data analysis, bringing together data from different equipment (Figure 7, 4th stage), and enabling more accurate management of inputs to optimize yield and connect information from different aspects of their operation that would be nearly impossible to do without this technology. We are on the cusp of the 5th stage shown in Figure 7, where all these individual systems can share information, allowing the whole farming ecosystem to evolve and unlocking even greater insights through the aggregation and analysis of huge data sets from a wide variety of sources.

The SPADE3 project's Web-based Asset and Vehicle data Exchange (WAVE) product team is focused on the abovementioned 5th stage. The scope of WAVE is to identify standardized systems for various agricultural data platforms to be able to efficiently communicate and share data. Currently most data platforms in the agriculture industry have different structures and processes for data sharing from one platform to another, or may entirely lack a data-sharing solution.

There are a few key aspects of data exchange that are out of scope for the project. Equipment manufacturers made it clear they need to be in control of direct communication with the machines they produce, for a variety of safety and reliability concerns. WAVE consequently focused on enabling data transfer among servers once data reaches the cloud; how the data reaches the cloud from the equipment is out of scope. Conformance testing and centralized servers / systems to enable data exchange are also out of scope.

The WAVE team's main focus is to identify existing standards for the communication of data to/from both fixed and mobile agricultural assets such as tractors, combines, grain dryers, irrigation systems, in-field sensors and related connected machines. Currently WAVE is working on eight different user-story-mediated areas: General Telemetry Data, Dashboard & Logistics, Sub-Field Operational Costs, In-Field Sensors, Third-Party Data Logging Systems, Radio-Frequency Identification (RFID) and Semi-Fixed and Fixed Assets, including grain storage and handling as well as irrigation systems.

As of this writing, the WAVE team has completed the user stories of interest. It is actively working with other SPADE and PAIL product groups to create process models and begin to capture data requirements. The team has also reached out to other organizations to identify potential standards:

- The Association of Equipment Management Professionals (AEMP) developed a construction-industry-centric equipment telematics API standard that could be adapted to agriculture equipment. This work is currently under review as an ISO technical specification (ISO 15143).

- The Agricultural Industry Electronics Foundation (AEF, http://www.aef-online.org) has an ongoing Extended Farm Data Interface (EFDI) effort to extend / share ISO 11783 ISOXML data for in-field telematics. The WAVE and EFDI teams are exploring collaboration opportunities.

The WAVE team is also coordinating efforts with PAIL and several SPADE project teams in evaluating the ISO 19156 (ISO, 2011) standard on Observations and Measurements for use with in-field sensors as well as irrigation systems.

OK to Spray

Safe operation is a central concept in responsible crop production. Specifically, minimization of spray drift has been increasingly relevant in terms of its regulatory implications for growers. Breaking the concept of safe operation down into processes and supporting data was a key deliverable of the SPADE Project's Crop Protection Product Team.

A conceptual contribution emerging from that team was the idea of "OK to Spray" (Figure 8): a process where actors in the application of crop protection products on a field evaluate, repeatedly as needed during application planning, preparation and execution, whether the necessary conditions for the operation are met. This concept was later generalized to "OK to Apply" and "OK to Proceed"

The goal is to reduce the chances of unintended non-compliance by providing a framework to make it easier for growers and other industry participants to communicate and understand the necessary conditions for a field operation to happen, as affected by the regulatory and commercial constraints to which the grower is subject.

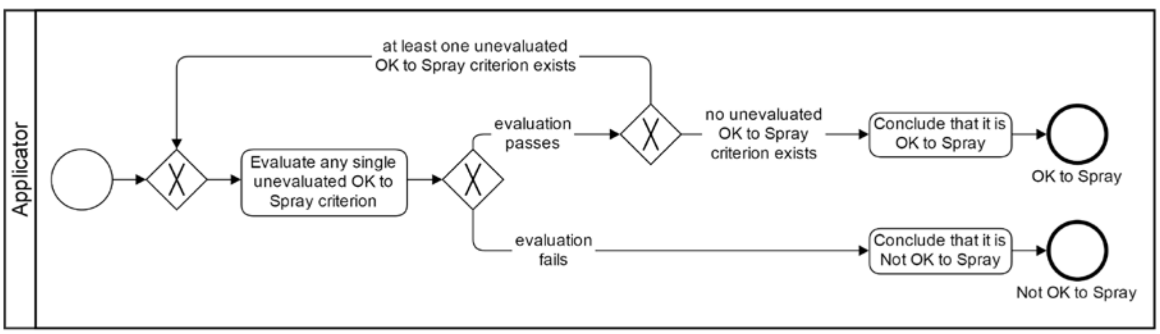

Figure 8: The "OK to Spray" process.

ADAPT: An Open-Source Application Programming Framework for Ag Data

One of SPADE's teams explored the feasibility of making an open-source tool to provide the industry:

- A common object model for field operations, and
- A plug-in-based framework for performing conversions between the common object model and proprietary formats. (Where a "plug-in" is defined as an external library, loaded at runtime, that contains the conversion logic.)

A proof-of-concept called "SPADE Conversion Toolbox" was deemed a success, and a new group was subsequently created within AgGateway to manage a new, follow-up effort: ADAPT. This group was established as a standing committee within AgGateway in order to provide long-term stability in the context of finite-length projects.

Figure 9 illustrates the operation of ADAPT with two high-level examples:

1. Incoming data from a Machine and Implement Control System (MICS, i.e., the controller in the cab. Data flow shown in thick red lines): A proprietary-format data file coming from a controller in the field is converted by a manufacturer-specific plug-in into an instance of the object model; a Farm Management Information System ("FMIS A") consumes the data.

2. Communication among FMIS (Data flow shown in thinner blue lines): FMIS A creates an instance of the object model, populates it with the data it wants to transmit, and uses the ADAPT plug-in to serialize it to a file. This ADAPT-formatted file is transmitted to another FMIS using the Internet or another means. (File transport is out of scope of ADAPT, accommodating different solutions available in the industry.) FMIS B uses the ADAPT plug-in to convert the ADAPT-formatted file to an instance of the object model, and then consumes the data.

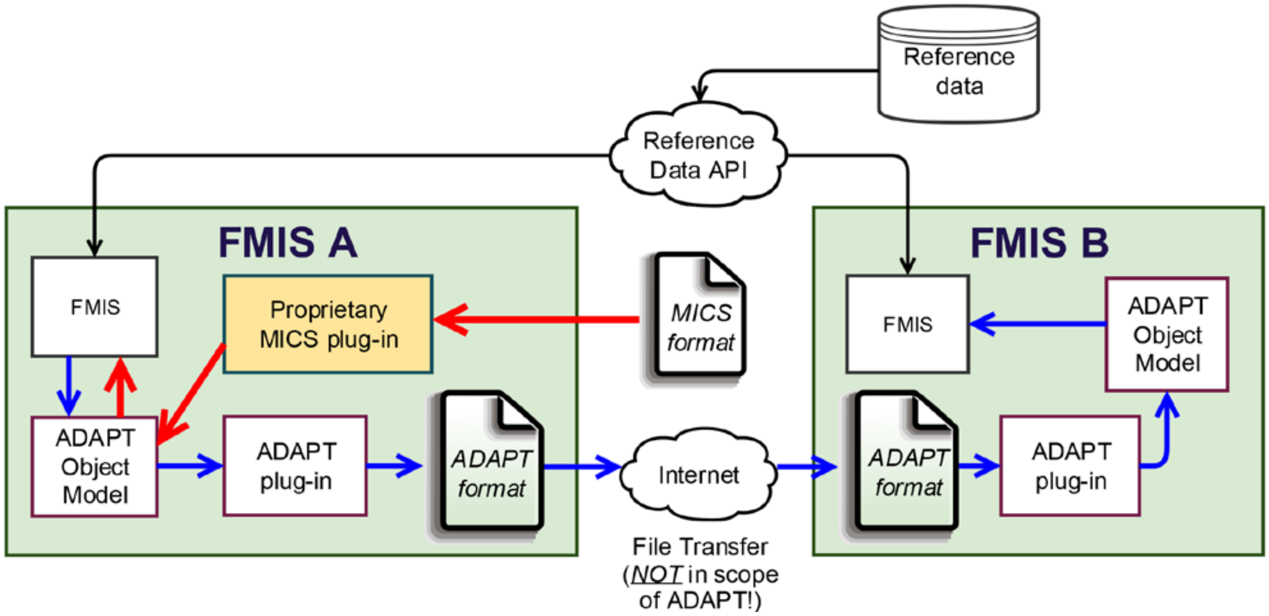

Figure 9: ADAPT examples, showing two farm management information systems (FMIS A and FMIS B) exchanging data.

Note how FMIS A and FMIS B are both supported by Reference Data.

The ADAPT team was diligent in exploring options for the licensing of deliverables, settling on:

- The open-source Eclipse Public License (Eclipse, 2004) for the common object model, the plug-in management tools, and some community-supported (e.g., basic serialization) plug-ins.
- Allowing individual plug-in writers to license their individual libraries as they choose.

These choices balanced a set of requirements, including supporting various manufacturer and service provider business/commercialization models, respecting companies' investments in patent portfolios, and promoting broad participation from both established industry players and startups alike, while providing a governance process that would seek to ward off unhelpful forks in code but enable participants to create derivative products. See more details at www.adaptframework.org.

Early in the development of SPADE and ADAPT tension developed between the need to represent data relevant to growers' business processes (e.g., US EPA numbers for products, Farm Service Agency Farm Numbers, etc.) and the desire for the definition of the Core Documents to be generic and portable across a worldwide industry. This led to the development of the ContextItem system.

Representing Geopolitical-Context-Dependent Data: a Balancing Act

Accommodating geopolitical-context-dependent data (GPCD; e.g., the mentioned EPA, FSA numbers) is important for this work to be relevant to a grower, but conflicting requirements must be reconciled:

- ADAPT developers must seek universality, staying free of regionally-specific clutter.
- However, different geographies' business processes involve context-specific data (e.g., EPA numbers). If these "context items" are not accommodated, the common object model's ability to represent the growers' business process, and therefore its relevance, suffers.

Additionally,

- It is desirable to use a controlled vocabulary. (To enforce shared meanings among the different actors exchanging information), However,
- the dynamic nature of business / regulation requires the vocabulary to be easily extensible.

ADAPT reconciled the contradictions by defining an object class, the ContextItem, that can be attached to various other objects in the common object model. A ContextItem is a key/value structure where the "key" code references a *ContextItemDefinition* that defines what each ContextItem means. The "value" is composed of a string value along with data needed to interpret it (such as a unit of measure) or a nested list of other ContextItems (e.g. PLSS cadastral information.)

AgGateway's SPADE project Reference Data API Product Team implemented a RESTful API to provide a machine-readable vocabulary of ContextItemDefinitions; its Standards & Guidelines Committee created an *ad-hoc* group to manage the vocabulary.

The ContextItem system can be used jointly with ISOXML's feature of associating unique IDs to its own locally-scoped IDs (defined in ISO 11783-10 Annex E.) This enables adding GPCD to ISOXML's otherwise generic and highly machine-specific scope, with no modifications.

Discussion

The body of work described herein results from the collaboration of dozens of industry experts over the course of five years. The SPADE, PAIL and ADAPT efforts have covered a lot of ground, so this paper is limited to providing a high-level overview. Some noteworthy emergent themes follow:

Expediency vs Consensus

One of the characteristics of the SPADE and PAIL projects was the expectation of expediency, which often seemed at odds with the need for a shared understanding, shared commitment and consensus among the participants. This was compounded by frequent turnover in the various product teams as individual participants' job responsibilities and assignments changed within their own companies.

When balancing expediency and consensus, the groups usually leaned towards consensus, which was seen as leading to durable support among project participants in the face of changing job assignments.

The Critical Role of Modeling Tools

Product-development teams found the three modeling tools—user stories, use cases, and BPMN diagrams—complementary and critically important. The user stories provided an unfettered, simple way to capture domain knowledge. Formal use cases extracted the essence of related user stories to document business processes in words. BPMN diagrams captured, in ways that were difficult (or impossible) with other tools, the elaborate business processes (having complex conditional logic and multiple parallel flows) that are quite common in field operations. The roles of these tools emerged along the way, as traditional use cases proved inadequate for some of the tasks at hand.

The Value of the ContextItem System's Data-Driven Approach

Early in the development of the Core Documents and ADAPT it became clear that adding geopolitical-context-dependent data to the Core Documents' data model was not going to be scalable or sustainable: it could only lead to a bloated model plagued with version-management issues. Two very important properties emerge from the ContextItem system's design:

- On one hand, it is a technical solution that enables the conveyance of sophisticated semantics, and enables the entry of ContextItems by FMIS that may have had no prior exposure to them (Daggett et al, 2016).
- On the other, it also offers relief to corporate IT departments that must necessarily commit to supporting specific versions of a standard for a prolonged period of time. By making the ContextItem system data-driven so new ContextItemDefinitions can be added to the API without having to change the undying object model version, organizations can add (and support) new business-process-specific data without incurring expensive version migrations.

Impact on Standards Processes

AgGateway's emphasis on collaboration, standards implementation (as opposed to standards creation) and active engagement of experts from both the technical and business sides of agriculture has led to very productive working relationships with other standards-related organizations, such as AEF, ASABE and various ISO subcommittees. This collaboration has addressed issues ranging from the AgGlossary (www.agglossary.org), to the need for unique identifiers in field operations documentation, to the potential uses and representations of OK to Spray. This level of collaboration affords a greater level of international traction to all the stakeholders involved.

Infrastructure-Building

Much effort went into architecting a distributed system for providing Reference Data to the industry. Multiple use cases had to be accommodated without favoring any particular business model:

- Some companies are interested in sourcing Reference Data for their own products.
- Other companies already have agreements in place to source their data through third-parties,
- There are sets of data that are best suited to be sourced through AgGateway itself (e.g., the ContextItem list, or data from companies that seek to tightly link their farm management-specific, Reference Data API-mediated data with their supply-chain-specific, AGIIS-mediated data.)

These requirements were reconciled through the design of a Reference Data Directory API that can enable an FMIS (or MICS) to find all available (registered) Reference Data sources. This mechanism, analogous to a phone book, is intended to simplify the discovery of Reference Data sources, as well as provide both third-party providers and own-data-sourcing manufacturers with an expanded, standardized market for their Reference Data products.

Future Direction: Implementation

AgGateway is an organization focused primarily on implementation. Consequently, the work done to date will only be complete once industry partners are using the data standards and infrastructure emerging from work in SPADE, PAIL and ADAPT to actually exchange Core Documents and Reference Data. Short-term future work will close some final gaps needed for this exchange to take place, such as adopting industry-standard best practices for data exchange authorization, authentication and supporting data services; for example, leveraging existing work done on implementation of OAuth 2.0 and OpenID Connect by the Open Ag Technology and Systems Group at Purdue University.

Conclusions

Strong market forces are demanding the rapid deployment of standard data formats for precision agriculture. The sheer need to feed an ever-increasing global population with limited resources in the context of an ever-increasing regulatory burden requires principled decision-making regarding crop input usage. This includes the capture and sharing of a variety of data points across multiple hardware and software platforms; i.e., greater interoperability.

AgGateway created its Precision Ag Council and subsequently chartered the SPADE (including its grain- and telematics-specific components CART and WAVE) and PAIL projects described above to address the need for greater interoperability. These teams followed a development process that fostered collaboration among a large group of domain and technical experts, including the capture of user stories, use cases, process models, data requirements; this was followed by the gap-checking (and proposal of extensions thereof, or new solutions, where appropriate) of existing data standards.

A sizeable body of work went into the development of infrastructure to support the desired outcomes: the ADAPT framework (www.adaptframework.org), the AgGateway standing committee created to house it for the future, and the SPADE Reference Data API Team's work on implementing proof-of-concept application programming interfaces.

Future direction includes further infrastructure development along the lines of adopting industry standard best practices for data exchange authorization, authentication and supporting data services, to be followed by implementation of Reference Data APIs and Core Document exchange.

Acknowledgements

The authors gratefully acknowledge the contributions of the dozens of other industry participants and AgGateway staff who participated in the work presented herein. Special thanks go to AgGateway member companies for their huge contribution of resources necessary for making the idea of eAgriculture a reality.

References

AgGateway (2015) The Ag Industry Identification System (AGIIS): A Critical Building Block to Facilitate eBusiness. http://s3.amazonaws.com/aggateway_public/AgGatewayWeb/About%20Us/AgGatewayAGIISFlyer_V_3_10815.pdf Accessed 9 June 2016.

Agile Alliance (2013) What is Agile Software Development? https://www.agilealliance.org/agile101/what-is-agile/ Accessed 9 Jun 2016.

Agile Alliance (2016) User Stories. https://www.agilealliance.org/glossary/user-stories/ Accessed 9 June 2016.

Daggett, D.G., Ferreyra, R.A., Reddy, L.T., Rhea, S.T., and Tevis, J.W. (2016). Filling in the blanks with ContextItems: a lightweight method for extending field operations object models. ASABE Paper No.2462418. ASABE, St. Joseph, MI..

Eclipse Foundation (2004) Eclipse Public License v1.0. http://www.eclipse.org/legal/epl-v10.html Accessed 9 Jun 2016.

Fielding, Roy Thomas (2000). "Chapter 5: Representational State Transfer (REST)". Architectural Styles and the Design of Network-based Software Architectures (Ph.D.). University of California, Irvine. http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm Accessed 9 June 2016.

GS1 (n.d.) Global Trade Item Number (GTIN) http://www.gs1.org/gtin Accessed 9 Jun 2016.

International Organization for Standardization. (2009) ISO 22006 Quality management systems – Guidelines for the application of ISO 9001:2008 to crop production. International Organization for Standardization, Geneva, Switzerland.

International Organization for Standardization. (2011). ISO 19156:2011 Geographic information -- Observations and measurements. International Organization for Standardization, Geneva, Switzerland.

International Organization for Standardization. (2015). ISO 11783-10:2015 Tractors and machinery for agriculture and forestry -- Serial control and communications data network -- Part 10: Task controller and management information system data interchange. International Organization for Standardization, Geneva, Switzerland.

Jacobson, I., Christerson, M., Jonsson P., Övergaard, G. (1992) Object-Oriented Software Engineering: A Use Case Driven Approach, Addison-Wesley.

Leach, P. J., Mealling, M., and Salz, R. (2005). “RFC 4122 - A Universally Unique IDentifier (UUID) URN Namespace.” IETF Tools. https://tools.ietf.org/html/rfc4122 Accessed 9 June 2016.

Object Management Group (2011). Business Process Model and Notation (BPMN) Version 2.0, http://www.omg.org/spec/BPMN/2.0/PDF/

Porter, M.E. and Heppelmann, J.E. (2014) How Smart, Connected Products Are Transforming Competition. Harvard Business Review, November 2014..https://hbr.org/2014/11/how-smart-connected-products-are-transforming-competition Accessed 9 June 2016.