CyNam
Agri-Tech Discovery

Data & Threat Model

Configured Things Ltd.
March 2023

Document Control

Version	Date	Notes
v0.1	27/07/2022	Initial Draft
v0.2	01/08/2022	Updated following feedback
v1.0	15/08/2022	Updated following feedback
v1.1	16/08/2022	Minor corrections
v1.2	18/10/2022	Added Data Aggregator Threat Model
v1.3	24/11/2022	Added Farm Threat Model
v1.4	06/01/2023	Added Trader / Retailer
v1.5	09/01/2023	Simplify Data Data Producer / Aggregator
v1.6	19/01/2023	Add details on approach and observations
V1.7	17/03/2023	Add exec summary and CAE sections

Contents

1 Introduction

Agri-Tech is an emerging area of interest within the cyber security sector, and is starting to transform UK food systems. With 78% of farmers across the UK using some form of Agri-Tech, it's vital to understand the security and risk components that lie within these technologies. The 2022 Government Food Strategy identifies innovation as a key component in meeting its economic and sustainably objectives.

Agriculture and food production is a major sector in Gloucestershire, and a key part of the regional economy. It is also the home of two leading Universities specialising in the future of agriculture, Hartpury University and the Royal Agricultural University (RAU), the selected site for what is currently described as the world's largest vertical farm[8], and the largest concentration of cyber technology businesses in the country.

Inevitably increased digitisation brings with it a corresponding cyber security threat. CyNam (Cyber Cheltenham) has therefore created a discovery project to bring together stakeholders from the Agri-Tech and Cyber communities to better understand the challenges and opportunities in this space.

1.1 Approach

Agriculture is a complex and specialised industry, with a diverse range of stakeholders. As such trying to understand how data flows through this ecosystem, and where the chains of influence are, is a significant barrier to entry; Before we could bring Cyber and Agri-Tech together we had to first create a shared representation of the problem space that both sides can recognise and relate to. The data model in this document is an attempt to do that, and has two aims:

- To provide a common baseline understanding across stakeholders in the Agriculture, Agri-Tech, and Cyber communities.
- To provide a framework against which we can quantify and ensure widespread representation in discussions

Although Agri-Tech is generally targeted at the farm, we wanted to also consider the ecosystem around the farm. Existing descriptions of the Agriculture ecosystem are based around the food supply chain. While relevant these don't necessarily provide a good basis for investigating the cyber specific aspects, and so we have attempted instead to create a description based around data; Where is it created, where is it held/ processed, and where is it consumed.

We then used the model to both identify groups it would be interesting to engage with, and as the basis for a discussion. These discussions were a mix of workshops bringing multiple stake holders together and some more focused interviews. In both cases we used the data model to draw out how data is used and what trust relationships are associated with it. We then used the STRIDE framework at an abstract level to categorise the potential threats, which also provided the non-cyber participants with an introduction to threat modelling.

The Data Model is described in section 2 and the Threat Models in section 3.

1.2 Executive Summary

Assessments of threats to Food Supply as one of the areas of Critical National Infrastructure consider the supply chain as a whole, and identify the current threats as ransomware (common to many areas) and the dependencies on other CNI areas, such as energy and

transportation. Individual producers are considered to be sufficiently diverse and numerous to not be a significant consideration. Because of this Agri-Tech, which is largely targeted at the producers, does not currently get much attention from a security perspective. In fact the key challenge is there there in no stakeholder to create a demand for cyber security; Adopting the technology is a business necessity rather than a choice for many farmers, who are in any case not positioned to make technical risk assessments. With security is not seen as a compelling differentiator there is little incentive for the the Agri-Tech companies, many of who are start-ups, to invest in that area ahead of functionality.

While similar security risks exist in other areas that are also rapidly adopting IoT and other forms of automation, such as smart cities and connected places, farms present a specific risk there is unlikely to be anyone with specific IT and security responsibility, compounded by the average of a UK farmer now approaching 60.

So within agriculture we can see many of the ingredients for a "perfect storm" on the horizon; a rapid expansion in the use of complex technology, driven by primarily by necessity rather than choice, and a limited ability or demand to consider the security implications, whilst the standardisation that comes from automation may erode the diversity that currently provides resilience.

1.3 Observations

The following is a summary of the observations from the discovery project.

1.3.1 Complexity

The diverse and complex of the Agriculture ecosystem makes it hard to model in a generalised way without continually acknowledging exceptions. We tried to address this with a generalised data model but then using specific case studies for the threat models. While diversity is generally good for resilience, complexity is rarely helpful for security.

1.3.2 Current Threat Perception

Food is categorised as one of the thirteen sectors of Critical National Infrastructure (CNI), although at this level it has a very broad perspective that includes factors such as global food supply, household-level food availability, and food safety.

Although by no means complacent, the general perception of the threat to Food is less than other CNI sectors. This is most recently summarised in the 2021 Food Security Report from DEFRA[4] which notes that in relation to the food supply chain in general:

> *"The UK is resilient to potential shocks in the food supply chain. Supply systems, which are owned and operated by the private sector, are adaptable and flexible in responding to problems. Government monitors risks and works with industry to respond to emerging issues and maintain supply chains.*
>
> *Notable risks to the supply chain stem from its dependence upon other critical sectors including energy, transportation, borders, labour, key inputs (chemicals, additives and ingredients), and data communications. In addition, the threat of cyber-attack to UK businesses, including those in the agri-food sector, is significant".*

Diversity is mentioned in several contexts as being a key factor in providing resilience to the food supply chain. The report does also note that:

> "*The UK's food supply chain is a highly complex system"*

and whilst that complexity undoubtedly supports the diversity, complexity and security are often in conflict within any system.

Specifically on cyber threats the report notes:

> *"The risk of cyber-attack to UK businesses is significant and continues to grow. It presents a threat to all CNI sectors. The nature of cyber-attacks means that they are varied and that attackers can adapt their approaches to their targets."*

So overall we can conclude that currently as a whole the Food sector of CNI is resilient, and that while it is exposed to the same growing cyber threats as other sectors, its diversity is seen as a significant mitigation.

Of course being part of the diversity doesn't help an individual producer or processor if they are attacked, so it is also important than individual participants in the supply chain remain independently resilient.

1.3.3 Emerging Threats

There are rapid changes in the way food is produced and distributed which are creating new attack vectors and threats. Market forces are driving producers to greater levels of automation in order to reduce costs and increase productivity, or in some cases simply to remain viable.

Both the on farm automation and the increased data integration throughout the food chain fundamentally depends on networked information systems, creating new and often unappreciated cyber attack vectors. Sensor systems designed and installed as IoT information sources become though automation control systems which require the rigour of industrial design. Increased dependence on such systems and between systems undermines the resilience which currently provides risk mitigation.

Automated systems remove diversity across farms, making many farms vulnerable to a common attack vector and therefore both a more attractive target and further undermining residence.

In Agriculture the latency between an automated decision being taken and the effect being seen in a crop is much longer than most industrial processes, so quick detection and response is of limited value in this context.

Because farms typically operate as price-takers, and at a relatively small scale compared to other parts of the food chain they are poorly equipped to understand and manage the cyber risks created by the systems the increasingly rely on; Remote management from vendors introduces its own set of risks to both the farm and the other systems it connects to.

1.3.4 Lack of Demand

While the benefits of cyber security in Agri-Tech are generally accepted it's less easy to see where the demand will come from.

Data and digital security is not at the top of a farmers list of worries; often they are having to adopt technology solutions to simply to remain viable. As a general point we got better traction when positioning the discussion as one of resilience (how do you keep your farm operating) rather that security (which seems to more linked to risks such as data theft)

None of the Agri-Tech startups we talked to identified security as a factor that is driving investment.

Cyber insurance in this space is generally immature and has a low take up. Underwriters in general do not seem to know yet how to factor security into lowering premiums.

The change to carbon and bio-diversity based subsidies might generate a need for data integrity, but in general the audit based approach seems to be considered adequate.

1.4 Further reading

For a more detailed background the following documents are recommended reading:

Cyber Security in UK Agriculture[1]

> Published in 2019 this white paper, a collaborative effort between Harper Adams University and NCC Group, addresses the cyber security threat to agriculture and the wider food network.

Cyber Risk and Security Implications in Smart Agriculture and Food Systems[2]

> Published in 2019 this paper from the University of Wisconsin-Madison provides a very good description of the new threats created by smart Agriculture.

Cybersecurity for the Internet of Things and Artificial Intelligence in the AgriTech Sector[3]

> Published in 2021 by the PETRAS National Centre for Excellence for IoT Systems Cyber Security this briefing provides an introduction to some of the threats and challenges.

2 Data Model

The following model is intended to provide a broad categorisation of the entities that exchange and process data within Agri-Tech. It is far from a perfect model, and many organisations may span or occupy multiple roles within it. Its main purpose is to provide a framework for discussion and a common language across stakeholders in Agri, Tech and Cyber.

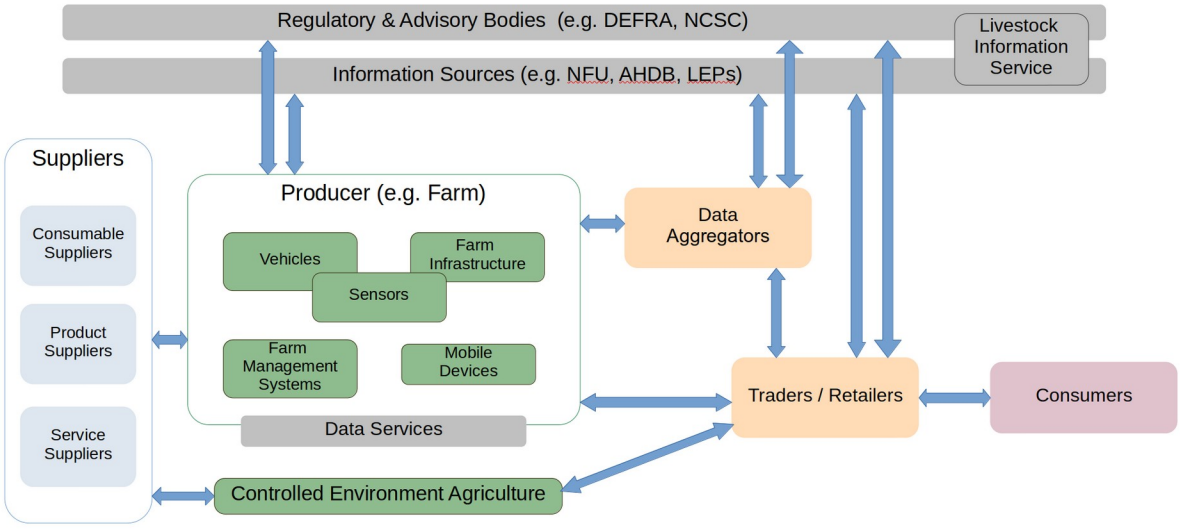

2.1 Regulatory & Advisory Bodies

This includes Government Organisations such as DEFRA (and their agencies such as the Rural Payment Agency), NCSC, etc

They are a source of policy and information data (e.g. DEFRA Data Services Platform[10]) that can be used to influence decisions.

There are requirements to provide some of them with data for payments (Grants, Sustainability Incentives, etc[12]) and statutory requirements such as livestock movements.

The NCSC (National Cyber Security Centre) have, in partnership with the NFU, published specific guidance on cyber security of farmers[6], covering the basic topics such updates, password management, antivirus, etc. The more generalised “Information for SMOs”[7] is also highly relevant here.

Note we may need to further distinguish here between statutory and non-statutory data.

2.2 Information Sources

This group captures the other information sources that feed into the Farm and Supply change. For example NGOs such NFU and LEPs act an industry trade body, influencing policy, providing information and advice, etc. Services such as the Met Office provide data which Farms may need to act on.

Note that the boundary between these two categories is blurred when it comes to trying to position individual stakeholders – for example the new Livestock Information Program[11] which consolidates a number of systems previously run by DEFRA is to be run by a limited company (Livestock Information Ltd) which is jointly owned by DEFRA and the AHDB, with

DEFRA funding the statutory elements and AHDB the value-add products. Likewise Agrimetrics (one of the four Innovate UK funded Agri-Tech centres funded) provides both commercial data sets and hosts the Defra Data Servcies Plateorm).

NFU and LEPs might fit into either or both categories.

2.3 Suppliers

For the purposes of the data model we split the suppliers into three groups based around the type of relationship they have with the farm:

Consumable Suppliers (e.g. Feed and Seed suppliers) have a short term transactional relationship where data is associated with each transaction. The data flow here is generally into the farm, and then passed on to the supply chain.

Product Suppliers (e.g. Machinery Dealers) have a longer term relationship linked to the lifecycle of the equipment, where the data relates to on-going support and maintenance. Data flows here are associated with the correct operation of equipment within the farm.

Service Suppliers (e.g. Veterinary, Agronomist, Agricultural Contactors) have a relationship with the farm that is a characterised by a two party contract (Farm and Service Suppler), as distinct from the more general Data Processors who are collecting data to be explicitly shared with some third party. In this case data is generally collected from the farm to processed externally and the results fed back to the farm.

The distinction between Product and Service supplier is the extent to which their solution requires on data being taken off the farm for processing. A Product supplier may for example collect data from the farm for fault diagnostics, and supply updates to product. A Service supplier will generally collect data, process / transform it in some way, and return a different set of derived data to the Farm.

2.4 Producer

The concept of producer is probably one of the most diverse within this model. The term Producer covers both Agriculture and Horticulture, and might represent a Farm or a Grower. We considered various forms characterisation such as size, tech maturity, type (arable, ruminant, etc) – but for the purposes of data flow and security it is most useful to decompose it into the broad types of things that generate, store, and used data.

We separate vehicles from other types of farm infrastructure (Environmental Controls (HVAC), Feeding, Milking) because their mobility generates a different threat profile.

Sensors includes anything that provides an automated source of data. They might be fixed (Soil / Water quality, door monitors, etc) or mobile (Livestock monitors, equipment tags, etc). There is an obvious overlap with vehicles and infrastructure that may include sensors; for our purposes we will make a distinction between sensors that are independent from the thing they are monitoring and those that are built in (controlled by and report to the thing they are monitoring).

Mobile devices (e.g. phones. Tablets, etc) are included as they are an increasing important interface and have a specific set of risks and connectivity concerns.

Farm Management System is a generic term for what may in practice be a collection of systems used to manage the farm. This area may need further decomposition.

2.5 Controlled Environment Agriculture (CEA)

Controlled Environment Agriculture (CEA), sometimes referred to as Vertical Farming, is a particular form of Controlled Environment Agriculture (CEA), where the process of growing food or other agricultural products is in a factory-style situation without the typical natural resources associated with plant production, such as soil and sunlight. These resources are instead provided via the use of controlled lighting and nutrient delivery technologies.

These environments have a tightly integrated nature, and are sufficiently different from the traditional farm that we have included them as a separate item in the model in anticipation that they may have a different type or style of data exchange. Examples in this space range from custom developed sites builds[8] to container based systems such as those from Lettus Grow which provide a kind of “farm-in-a-box” approach.

2.6 Data Services

This component is included to capture and represent that even systems which appear to operate just within the context of a farm (such as sensors reporting into a Farm Management System) may have data paths that leave and return to the farm.

2.7 Data Aggregators

Farms use a number of services that collect data and make it available to a number of third parties; Examples include Carbon Audit companies, Drone Crop Mapping, etc.

Data Aggregators address the need to connect data providers and consumers across the food supply chain. They provide value to both data provides and consumers by acting as a single point of contact and managing the required data aggregation and transformation services.

The distinction between Farm Data Services and Data Processors is who the data is shared with; This may be an artificial split and in practice all data that leave the Farm should for a cyber perspective be considered as “shared” with a third party, but initial discussions have suggested that this is a useful split to consider.

2.8 Traders / Retailers

This area encompasses an extensive and complex network of Traders, Processors, Wholesalers, Retailers and Outlets, which eventually make food available to Consumers. While it may seem like a vast over simplification to reduce this to single role within the model we justify that by observing that this is the most mature part of the eco-system and parallels other supply chains in terms of considerations and threats. It is also one of the three areas identified for a more detailed study.

2.9 Consumers

Consumers are of course an obvious endpoint of the Food Supply chain. In the context of the data model consider them from the context of both how they can trust the data / labelling on food and its origins, and how their demand for that data might affect the end to end processing of that data.

3 Threat Models

The following sections provide threat models for various parts of the data models, produced by through a combination of workshops and meetings with representative stakeholders.

Note: The threats captured here should be considered as a way to better understand the attack surface of a part of the data model – they do not imply that they exist or are not mitigated in specific systems.

For each part of the system we create an abstracted model using a standard notation, focusing on interfaces, data stores, and trust boundaries. We keep these models as simple as possible, so that we can capture the essence of the system rather than a specific implementation.

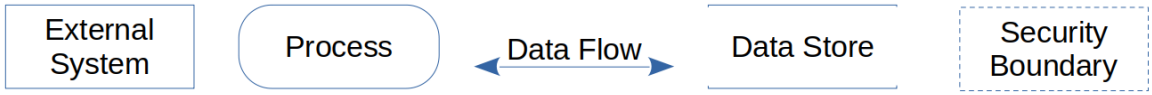

We then identify threats using the STRIDE framework[13], which uses the following headings to guide the analysis:

Threat	Definition	Desired property
Spoofing	Pretending to be someone else	Authenticity
Tampering	Changing Information	Integrity
Repudiation	Denying an act	Non-repudiability
Information disclosure	Revealing information to someone not authorised to receive it	Confidentiality
Denial of Service	Preventing users from accessing the service	Availability
Elevation of Privilege	Having access without authorisation	Authorisation

The STRIDE model is used because of its simplicity and applicability at various levels of abstraction.

Note that threats do not exist in isolation, and most interfaces by their very nature are exposed to all of the above to some extent; Spoofing an Identity can lead to Information disclosure, and also creates a repudiation risk. Elevation of Privilege, when an attacker not only claims to be a valid user but also one with extended authorisation, is often the most threatening area as it can build on all of the other threats. The aim of the threat analysis therefore is not to provide a comprehensive list, but to provide a context in which to think about the security risks to a system.

3.1 Producer

As noted above the concept of “Producer” in the data model covers a very diverse range of entities, and one which is evolving rapidly in the use digital technology. Therefore rather than attempt a generalised threat analysis for this whole we look at one particular subsystem in detail which provides good insight into the types of threat a producer may be exposed to.

3.1.1 LoRaWAN Sensor network

Sensor networks are a key aspect of digitally enabled farms, providing a continual stream of data from a range of sources, such as:

- Soil Sensors
- Drinking trough levels
- Tank levels and temperatures
- Environmental conditions for livestock housing
- Smart Rumen Bolus
- Building Protection

Although a number of sensor technologies exist, LoRaWAN is a good fit for agriculture use cases as it requires very little infrastructure; the sensors are battery powered, can be easily deployed, and have a range of several miles to a gateway which acts as a bridge between the radio network and an IT network. The protocol it uses is specificity designed for small amounts of data sent on an fairly infrequent basis (minutes or hours).

While the majority of devices are sensors (i.e. monitor / collect data and send it on a periodic basis) it is also possible to have LoRaWAN based activators, devices which listen for a control signal and perform some actions, such as opening or closing a valve. Listening generally consumes more power that a periodic transmission, especially if the device has to listen frequently to provide a low latency response, plus of course more power is required for mechanical activation. Activators also tend to be in a fixed location, and so don't need the mobility / battery-based characteristics of sensors.

For the purposes of the threat analysis we use the following architecture pattern.

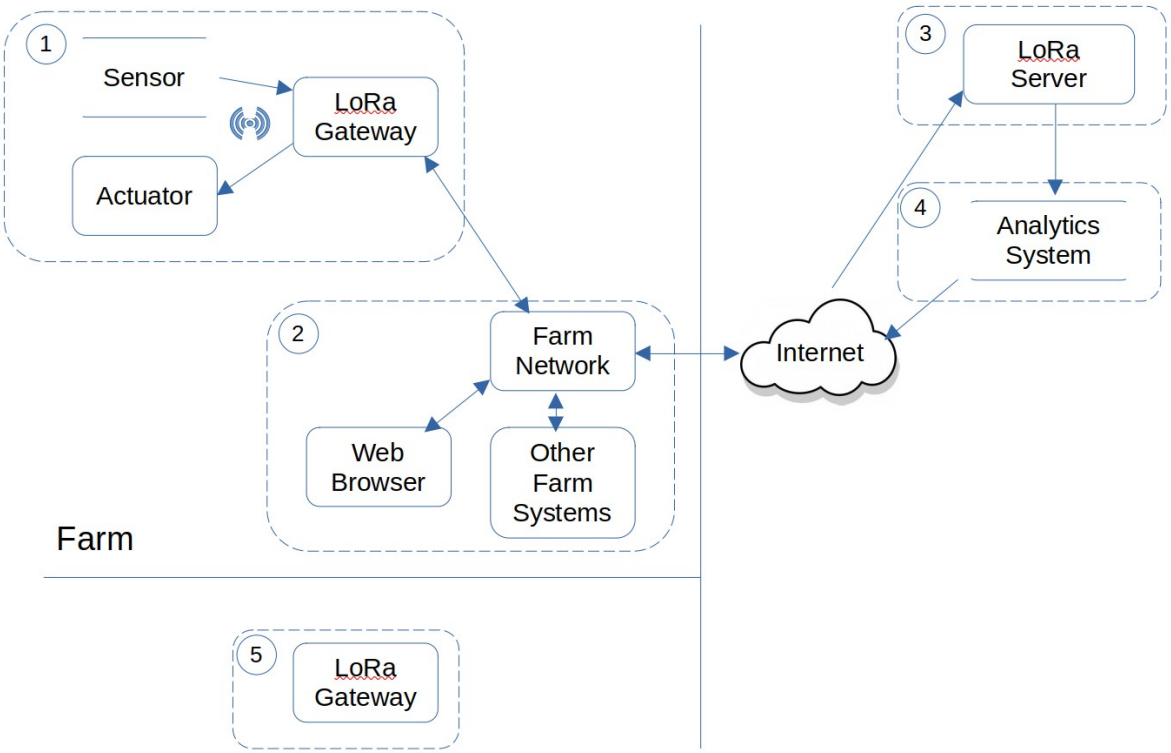

This Architecture has five trust domains.

1	The On-Farm LoRaWAN devices and gateways, provided and managed by a service provider.
2	The Farm network and other systems connected to it, owned and managed by the Farm. This assumes that the gateway is using this network as its backhaul to the LoRaWAN server.
3	The LoRaWAN server, typically provided as a cloud hosted service and shared with other customers.
4	An Analytics system providing visualisation of the data, and in some cases control signals derived from the data. Although this may be presented by the LoRaWAN provided as an integral part of the solution in reality it is often hosted and operated by a third party
5	A LoRaWAN gateway which is not part of the solution but is within range of the sensors. The design of LoRaWAN is such that any gateway within range receives all packets for a sensor.

It also has two Data Stores:

A sensor can act as a data store if it accumulates values to be sent – for example a sensor acting as a meter which sends a count build up over a period of time is acting as a store.

The analytics system provides the main store of collected data

3.1.1.1 LoRaWAN security characteristics

The three key components in a LoRaWAN network are the Devices (which send data over the radio network and listens for a response in a predefined time window), the Gateways (which listen for all radio traffic and forwards it over an IP network) and the LoRaWAN Server[1] (which decodes and decrypts the messages).

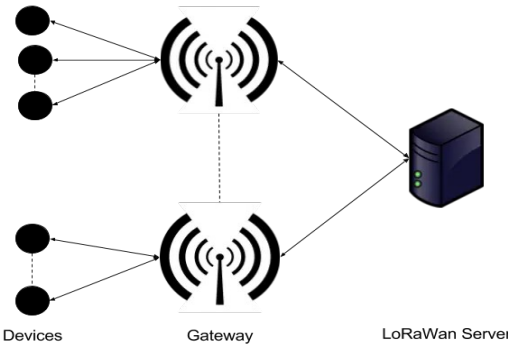

Devices must be within range of at least one Gateway (which depending on the physical geography can be several miles), whereas the distance between the Gateway and the network server is effectively unlimited; A common pattern is for the LoRaWAN Server to be hosted in the cloud.

In general the LoRaWAN security depends on the security and integrity of a shared secret between each Device and the LoRaWAN server. This is used, either directly or via derived session keys, to protect the integrity of the message and to encrypt the data payload. The

[1]Technically this is three services, Network Server, Join Server, and Application Server, but this is level of detail is irrelevant in this context, assuming they communicate via secure protocols.

secret itself is never sent over the network, so any intermediate device is unable to get it just by eavesdropping on the network. Provided this secret is well chosen and remains protected then in the context of this analysis be considered to be reasonably secure given the value of the data, especially if using Over the Air Activation (OTTA) with the V1.1 protocol. If however the secrets are made easy to access (e.g. readable from the device) or easy to guess (e.g. follow some pattern) then the network becomes vulnerable.

As a radio based protocol it is vulnerable to various Denial of Service threats.

Gateways are generally simple protocol convertors (LoRa to IP), forwarding all packets to the receiver. The simplest protocol between a gateway a LoRaWAN server is based on UDP, which is OK in so far as the data integrity and contents are protected by the LoRaWAN protocol, but this doesn't allow the network server to verify the gateway's identity which can expose it to a variety of attacks. Most gateways can be configured to support a more secure protocol such as Basics Station[14] or MQTTS which should always be used instead.

3.1.1.2 Threats

3.1.1.2.1 Spoofing

There are a number of places in the system where identity spoofing is a threat.

- A device's identity is protected by a shared secret being known only to the device and the LoRa server. If this secret is leaked then it is possible to spoof a device by restarting the join process. By their nature devices have a limited configuration capability, so there are number of ways in which this could happen:
  - In some devices the secret is added to the device during manufacture, and so has to be provided by various links in the supply chain – any of which could be compromised. Or it may simply be printed on the device casing.
  - Devices which are configured via an app on smart phone (e.g. NFC or WiFi) may also allow the configuration to read.
- Any IP interface (i.e. all of those apart from the Device / Gateway) are exposed to spoofing if secure protocols are not used
- Where a component supports remote access, for example remote support of a Gateway, then identity spoofing presents a specific risk to any network the device has access to.

3.1.1.2.2 Tampering

All network-based interfaces are exposed to the threat of tampering, but most can be mitigated through the use of secure protocols. The LoRaWAN protocol include mandatory security to reflect that its data is broadcast rather than directed to a specific target.

The exposed and sometimes unobserved location of devices and gateways makes them exposed to various tampering threats which could be as simple as changing the immediate environment of the device to create false readings.

Where data leaves the farm and is then returned via the analytics system there is limited or no protection on the integrity of the data, and this path crosses a number of trust domains. If there is a control loop between a sensor and actuator then the safety aspect of this needs careful consideration.

3.1.1.2.3 Repudiation

Where a LoRaWAN device is used to provide a general input (such a button to record that a location needs attention) then that is in effect shared input with no traceability. If the originator

is important (i.e. there is a need to know who is reporting) then individual portable devices can be used.

The limited capabilities of a LoRaWAN device make it impractical to establish an audit trail of actions at an actuator, and whilst the protocol requires knowledge of the shared secret there is no direct traceability to the origin to know if the secret has been compromised.

Any changes made by a service provider, in either the hosted part of the solution or remotely managed infrastructure depends on the operation processes of the provider to provide non-repudiation. It is in both parties' interests to ensure this is in place.

3.1.1.2.4 Information Disclosure

The design of LoRaWAN means that any gateway within range of the devices will receive all messages from them. While the data payload is encrypted this does still disclose information about activity in certain locations and metadata to the receiving gateway.

The most significant risks of disclosure are in the hosted services, where security depends on the operational process of the provider.

3.1.1.2.5 Denial of Service

The exposed and sometimes unobserved location of devices and gateways makes them vulnerable to physical damage and denial of service, as does harsh environmental conditions, often the physical protection of a device or gateway makes up much of its cost.

Devices have to minimise the time used to send and listen for data. Depending on the device configuration various threats exist which can force a device to retransmit data or to listen more frequently that intended, both of which will drain the battery and create a Denial of Service.

Rogue devices within range of the gateways could flood the network with invalid packets. Unless network filtering is used on the gateways all traffic has to be sent to the LoRa server, so all gateways within range are susceptible.

A compromised gateway could be used to launch a Denial of Service attack on the Farm network and other systems connected to it. This is a participially significant attack vector if the gateway provides some form of interactive remote access, e.g. for support.

In cases where a control loop exists between a sensor and activator any loss of service of the Farm network or its connection to the internet may create safety considerations.

3.1.1.2.6 Elevation of Privilege

The majority of privileged operations are in the configuration of the LoRa Server and Analytics system. Where these are based on a shared platform the threats are dependent on the operational processes of the service provider to provide isolation between customers.

Remotely Managed devices on the farm may create an elevation threat to other systems on the farm network, since they may be considered by default to be part of the same trust domain.

3.2 Data Aggregators

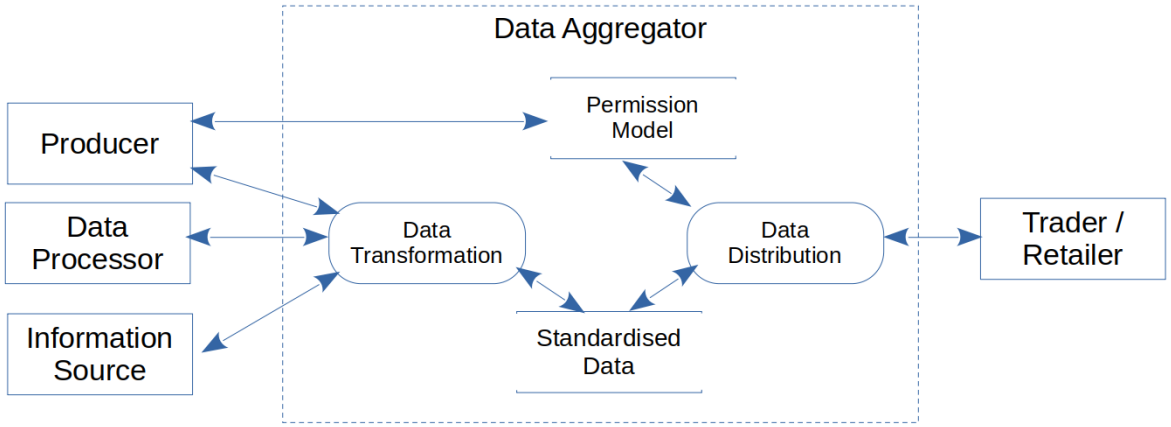

Data Aggregators collect data specific to a Producer, either directly or indirectly from Data Processors on their behalf. They may also collect data from publicly available sources. They transform it into a standardised form, and supply subsets of the data from a number of Producers to a Trader/Retailer. Each Trader/Retailer may have different requirements. The contract to supply the data is formed directly between the Consumer and Trader/Retailer; with the Date Aggregator acting as a broker. The data remains owner by the Producer, and is only provided with the specific consent of the Producer.

3.2.1 Key trust relationships

The Producer trusts that Data Aggregator will not change their data, nor release it without their permission.

The Producer is also trusting that the Data Aggregator is correctly representing the Trader/Retailer.

The Trader/Retailer is trusting that Data Aggregator provide it with accurate data.

3.2.2 Threats

3.2.2.1 Spoofing

There are a number of places in the system where identity spoofing is a threat.

- If the Data Aggregator fails to establish the identity of a Trader/Retailer it will misrepresent the permission it is asking from the producer.
- If the Data Aggregator fails to establish the identity of the Producer it is exposed to tampered input data, and may disclose information about the relationship between Producers and Trader/Retailers, and an elevation of privilege in accepting changes to the permission model.
- If the Producer or Data Processor fails to establish the identity of the Data Aggregator they may disclose information to a third party, including credentials needed to access data on their behalf from a Data Processor.
- If the Trader/Retailer fails to establish the identity of the Data Aggregator it may disclose information about its relationship with Producers, and be exposed to tampered input data.

3.2.2.2 Tampering

All network-based interfaces are exposed to the threat of tampering, but most can be mitigated through the use of secure protocols. Specific additional threats to consider include:

- If data is provided by the Produce in file format, then both the transport (e.g. e-mail) and the documents can be tampered with.
- Data from other Information Sources (such as weather data) may be coming from parties that are outside any specific contractual relationship, and is likely to have an impact across multiple producers. The risks derived from the use of this data therefore need specific consideration.
- Where the collected data is transformed into a standardised form it may be hard to validate the integrity of the data with the originator; even if made available to the Producer it may not be in a form that they recognise. The transformation is therefore itself exposed to a threat that may be hard to detect, and which has a significant impact.
- The Aggregator has to maintain record of what data each producer is willing to share. The integrity of these records is a key factor in the trust relationship between the producer and the Aggregator. Tampering with these would lead to further Information Disclosure threats.

3.2.2.3 Repudiation

The main repudiation threat is around changes to the permission model since this is the mechanism by which the Producer gives permission to the Data Aggregator to disclose specific data.

- If its possible for the Data Aggregator to repudiate what permissions it asked for, or the Producer to repudiate what permissions it granted then a major aspect of the trust model is exposed. Since the contact of what data needs to be supplied is between the Producer and the Retailer/Trader it is of course possible for Producer to verify that requests are consistent with the contact, although how practical this depends on the complexity of the data.
- If the permission to disclose is open ended (i.e. not for a specific period of time) then the repudiation threat also applies to the removal of permission; since the contact is between the Producer and Retailer/Trader the Data Aggregator there is a risk that that they will not be notified of this change.

3.2.2.4 Information Disclosure

A number Information disclosure threats have already been covered as a consequence of other threats, particularly spoofed identity. Additional threats exist in the following areas:

- All interfaces have the risk of disclosure if they do not use secure protocols.
- The aggregated data held by the by this system would make it an attractive target for attackers, so in addition to leakage though the interfaces the controls around the data stores need to be carefully considered, especially if the system is cloud hosted. This includes not only the data itself, but the metadata relating to permission (which exposes details of the relationships between Produces and Traders/Retailers), and any credentials held to access data from a Data Processor on behalf of a Producer.

3.2.2.5 Denial of Service

Clearly a major denial of service would affect a significant number of Producers and Trader/Retailers, although the impact would depend on the duration and timing.

- In most cases the data is not time critical, and so short interruptions in services, providing they do not result in data loss, can be tolerated.
- It was however noted that there are some periods related to harvest when the data does become time critical, so an attacker aiming for disruption may focus on specific dates.
- The seasonality of agriculture may also lead to high peak data loads that need to be accommodated to avoid this also becoming a Denial of Service threat.
- Given the complexity of the data aggregation its not clear if there would be an easy mitigation to a significant duration outage, for example from a ransomware attack, nor if responsibility for this should rest with the Data Aggregator or the Producer and Trader/Retailers who are ultimately the risk owners.

3.2.2.6 Elevation of Privilege

Assuming that all Producers and Trader/Retailers have separate identities and roles within the system there is no reason for any of them to need elevated privileges. The main elevation of privilege threat is therefore from accidental or malicious misconfiguration.

- Complexity within the system configuration may both increase the risk of misconfiguration and make it difficult to verify that configuration is correct.
- Social engineering attacks may lead to compromised credentials with elevated privileges.

3.3 Traders / Retailers

As noted in the introduction to the data model, this element encompasses an extensive and complex network of Traders, Processors, Wholesalers, Retailers and Outlets – basically all and any of the processing of food between the farm and the consumer. While that may seem like too big a simplification, the justification is that all businesses in this sector are industrial operations that have well established practices for data handling, resilience etc, with dedicated IT and security functions. Any deep investigation in here would be too complex for the scope of our work, and while threats undoubtedly exist they are unlikely to be specific to this domain.

We did talk to a representative of a major food processor to understand the nature of data transactions with a farmer, and the key learnings from that were as follows.

All produce, whether it comes direct from the producer, via a co-operative, or imported, etc has data that identifies its origin and history. If the data hasn't reached and been accepted by the processor by the time the good arrive then they simply won't be accepted. That data is then tracked against and augmented as the produce as it moves through the production line / supply chain.

The quality / integrity of the data that comes from a farm is normally assured via some kind of audit process on the farm (such as Red Tractor). That generally sets the bar that the farmer has to meet to satisfy the processor.

Demand data from the processor, telling the Farm what is required, is normally a long term contractual commitment rather that dynamic data.

Some initial thinking in this area suggested that the use of blockchain technologies could provide increased assurance of the integrity of the data, and bring the resilience benefits of a decentralized system. IBM, for example, offer a hosted blockchain based system[15]. Other studies have suggested that while this does provide data-management benefits, the key challenges are in the initial data capture.

Its worth noting that the data exchange could be via a Data Aggregator, which is covered elsewhere. Here were looking just at the case of the direct exchange of data related to a specific shipment.

3.3.1 Key trust relationships

The scope of any trust is fairly limited, since the data exchanges are in effect linked to specific transactions (goods moving from Farm to Producer). The main aspect of the trust is that it wouldn't be in either parties long term commercial interests to not supply accurate data.

3.3.2 Threats

3.3.2.1 Spoofing

While its possible that someone could spoof either the Farm or Processor, the absence or duplication of data would prevent the produce from being accepted.

3.3.2.2 Tampering

The accuracy of the data is based on long term auditing (e.g. Red Tractor Accreditation) and periodic sampling. For those aspects of the data used to assure that the produce are of a sufficient quality, tampering could have an effect in two directions; Reducing the apparent quality could create a DoS attack, and increasing the quality could direct produce into the wrong part of the food chain, but this would required contracts and processing that spanned multiple supply chains.

3.3.2.3 Repudiation

Nothing noted.

3.3.2.4 Information Disclosure

Nothing noted.

3.3.2.5 Denial of Service

Any missing or incorrect documents ion that prevents the produce from being accepted would create DoS attack, targeted at the Farmer but it is also easily detected given the direct relationship to physical produce. As noted elsewhere the main mitigation from a food chain security perspective is the diversity of suppliers; As produce is aggregated through this processing chain (several farmers sharing a co-operative, etc) the risk increases, but so does the size and capability of the organisations involves, which has a countering effect.

3.3.2.6 Elevation of Privilege

Nothing Noted.

3.4 Controlled Environment Agriculture (CEA)

The CEA approach is to maintain optimal growing conditions throughout the development of the crop by controlling or removing the the variability of outdoor environmental conditions. Production takes place within an enclosed growing structure which can vary in scale from a container to a plant factory, within which items such as light, temperature, humidity, CO_2 and nutrient levels are carefully controlled.

The potential benefits include an all year growing period optimised for yield and quality, reduces manual labour, reduced chemical consumption for weed and pest control, localised food production, and lower water usage.

Within this there is a spectrum of solutions from fully controlled enclosed, closed loop systems, to automated glasshouses. The former are closest in nature to the kind of industrial control systems sometimes described as “Industry 4.0”, where as the latter are an interesting mid point from more traditional production methods.

Although the fully enclosed approach is still very much in the “emerging tech” category, it seems to be an area of considerable interest to investors. From a cyber security perspective the self-contained nature offers more possibility for a strong, designed-in, enforcement layer than the situations when IoT is being added to existing facilities. However, perhaps reflecting that the IP in such systems sits within the effectiveness of the the control loop, very little public information is available on the architectures, with most published research focusing on the details of the environmental controls. In the timeframe available for this study we were unable to engage with any of the CAE based producers in our region.

As with any closed loop control system, the main threats will come from whatever limited connectivity is required to the outside world. Physical access is of course still needed to maintain the system and harvest the the crops, and provides one obvious attack vector. It seems that the control systems are finely tuned to specific crops, although presumably the system all use similar base systems and components, which creates a possibility of class attacks.

Some systems specifically offer remote management and monitoring from smart devices, which of course can introduce its own attack vector. At least one offers the ability to share data with an expert support team, which depending on the level of access given may create a trust relationship which goes beyond the of advisor, and create an additional attack surface.

The “sealed box” type isolation of these systems can be both a benefit and a risk; by being isolated the attack surface is reduced, but if it is compromised in some way then the issue may go undetected.

4 References

[1] CyberSecurity in UK Agriculture – NCC Group: https://research.nccgroup.com/wp-content/uploads/2020/07/agriculture-whitepaper-final-online.pdf

[2] Cyber Risk and Security Implications in Smart Agriculture and Food Systems: https://jahnresearchgroup.webhosting.cals.wisc.edu/wp-content/uploads/sites/223/2019/01/Agricultural-Cyber-Risk-and-Security.pdf

[3] Cybersecurity for the Internet of Things and Artificial Intelligence in the AgriTech Sector – PETRAS Industry Briefing

https://petras-iot.org/wp-content/uploads/2021/04/PETRAS_IndustryBriefing_Agritech.pdf

[4] United Kingdom Food Security Report 2021 https://www.gov.uk/government/statistics/united-kingdom-food-security-report-2021/united-kingdom-food-security-report-2021-theme-3-food-supply-chain-resilience

[5] An overview of PERA and the Purdue Methodology https://link.springer.com/content/pdf/10.1007/978-0-387-34941-1_8.pdf

[6] Cyber security for farmers – NCSC / NFU https://www.ncsc.gov.uk/guidance/cyber-security-for-farmers

[7] NCSC Information for Small & Medium sized organisations https://www.ncsc.gov.uk/section/information-for/small-medium-sized-organisations

[8] The world's largest vertical farm https://www.jonesfoodcompany.co.uk/our-latest-news/the-worlds-largest-vertical-farm

[9] Drop & Grow Container Farms: https://www.lettusgrow.com

[10] Defra Data Service Platform: https://environment.data.gov.uk

[11] Livestock Information Program: https://ahdb.org.uk/LIP

[12] New Farming Policies and Payments in England

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1093812/payments-for-farmers.pdf

[13] STRIDE threat modelling
https://en.wikipedia.org/wiki/STRIDE_(security)

[14] LoRa Basics Station
https://en.wikipedia.org/wiki/STRIDE_(security)

[15] IBM Supply Chain Intelligence Suite: Food Trust
https://www.ibm.com/products/supply-chain-intelligence-suite/food-trust