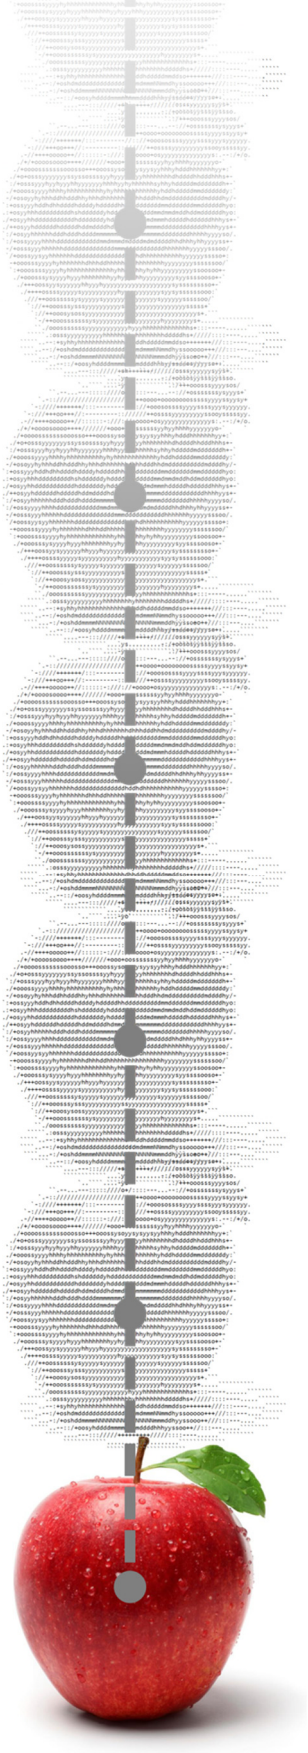

Blockchain's Role in the Produce Supply Chain

The buzz around blockchain and distributed ledgers has reached a fevered pitch,[1] with numerous applications[2] being bandied about including its use in supply chains. A lot has been written specifically about the potential use of blockchain technologies in fresh food supply chains. The vision is that industry-wide blockchains can provide:

- *Provenance—stronger assurance of origin and chain-of-custody*
- *Recalls—faster and more precise recalls*
- *Freshness—fresher produce and meat, reduce waste and spoilage*
- *Safety—fewer contamination incidents*

A strong case can be made that the greatest value potential is in these last two, improving freshness and safety. Nevertheless, most of the pilots and discussion to date have focused on the first two; how blockchain can help establish provenance and recall capabilities, in particular by providing trace-ability. To understand the role of blockchain in providing traceability (and hence provenance and recall), we examine four ways to achieve traceability:

- *1-up/1-back proprietary—The most common approach today, traceability data is held within each firm in a proprietary format*
- *1-up/1-back standards-based—Data is still stored in each firm, but in a standard way*
- *Centralized Networked SaaS—The best way to get end-to-end traceability, this emerging approach has a chain-wide database*
- *Decentralized Blockchain-based—Blockchain is getting lots of buzz but is most powerful when combined with a networked SaaS platform.*

Then we take a look at what is required for the four use cases (provenance, recalls, improving freshness, and safety) and blockchain's potential contribution to these, including the role of smart contracts.

[1] Bitcoin's recent astronomical bitcoin exchange rates have helped fuel the mania.

[2] Applications of blockchain being developed range from securities trade clearing and settlement to identity management, insurance, deeds registry, micropayments, smart grid/energy trading, crowd-funding, regulatory compliance, and much, much more. Many lists of blockchain applications and companies have been published such as: Business Insider, MEDICI, CB Insights, Blockgeeks, Blockchain Daily News (250 top blockchain companies).

Contents

Provenance and Recall Capabilities

The Desire for Better Provenance Assurance and Recall Capabilities

Provenance Assurance—Increasingly consumers, retailers, and brand owners want more precision and confidence in knowing where their food has come from. Consumers desire authenticity for food that claims to be grown in a particular region or following a particular regime (e.g. organic, grass fed, etc.). Some brand owners want to correlate the quality, robustness, and yield of produce with the specific location and date of harvest, in order to tease out farming and harvesting best practices and/or to help in identifying and breeding fruits and vegetables with desirable characteristics such as color, taste, shelf life, and yield.

Recall—Brand owners and retailers would like to identify any incidents of contamination at the earliest possible time, reliably and swiftly trace the contamination back to its source, reliably and swiftly trace forward to where all potentially contaminated produce has been sent, and execute high-confidence, rapid, precise recalls. Ideally, tainted produce or meat is caught before product has been put on the retail shelf.

Traceability's Central Role in Provenance and Recall Capabilities

Traceability is fundamental to achieving provenance assurance and robust recall capabilities. Upstream traceability (aka backwards traceability) is the ability to trace an item[4] back to its original source.

What is Blockchain

Blockchain is technology providing an immutable, secure, decentralized, shared ledger. Breaking down this definition (in reverse order):

- Ledger—Similar to a financial ledger, a blockchain is an ordered list of 'transactions.' Transactions can be financial exchanges (as on bitcoin and other crypto-currency blockchains) or the record of an event, such as a hand-off in a chain of custody, the completion of a task, or even a temperature reading at a specific time and place. The sequence of validated transactions is strictly maintained. For cryptocurrencies, this prevents double-spend, as the oldest transaction prevails and newer duplicates are invalidated.
- Shared—The ledger provides a common view (aka the single-version-of-the-truth) across many parties, even if they lack trust or a longstanding relationship.
- Decentralized—The ledger is broadly replicated and verified by a large number of different entities. This makes blockchains robust and resistant to cyberattacks and system failures, as an attacker or outage would have to impact a large portion of the network of entities in order to compromise the blockchain. The degree of decentralization varies between different blockchains, depending on the goals and design.[3]
- Immutable—The ledger is immutable, meaning that transactions can never by deleted or modified. If there is an error in a transaction, it can only be corrected by writing a new transaction undoing the error. This makes blockchains highly auditable, as there can be no after-the-fact rewriting without a clear trail of what happened.
- Secure—Cryptographic technology (e.g. hashes, digital signatures) is used to ensure there is no tampering with the data in the blockchain. The way in which hashes are used to link the content of each block of transactions from/to the previous and next blocks is how blockchain got its name.

[3] Bitcoin, Ethereum, Hyperledger, R3 Corda, Ripple, and numerous other blockchain technologies have been developed with different purposes and design goals in mind, and hence with very different capabilities and performance characteristics. Some of these technologies, such as Hyperledger, are a collection of various types of blockchains and tools. Hyperledger is an open source project which currently consists of five different frameworks, each designed to support different types of blockchain environments, goals, and constraints (e.g. minimal resource validators, digital identities, smart contracts, application development, and mobile applications).

[4] For manufactured items including processed foods, this includes tracing the constituent parts and materials back to their ultimate sources, often through multiple tiers of production and/or blending of batches. For fresh cut produce, the processing is generally simpler than for canned or packaged foods. Fresh cut processing may include trimming, cutting, washing, drying, and bagging of fresh salad vegetables, during which blending of different sources of produce typically occurs.

Downstream or forward traceability is the ability to start at the source or any intermediate point and find out where each end item from any particular batch ended up. A traceability system should provide both upstream/backwards and downstream/forward traceability.

The most widely used approach to traceability in the food supply chain is one-up/one-back traceability (aka 1-up/1-down), where each participant in the end-to-end chain (grower, processing plants, distribution centers, retailer, etc.) keeps track of where every item (or batch of items) they received came from, and where each item (or batch or lot) is sent to. One-up/one-down food traceability is mandated in the United States by Section 306 of the 2002 Bioterrorism Act.

If everyone in the chain does 1-up/1-down traceability record keeping properly, then when a problem occurs, such as a batch of food is contaminated, the problem can be traced back to the source of contamination anywhere in the chain. Contamination can happen at different points in the chain, not just at the origin. When contaminated food is discovered (either by testing or when people get sick), backwards traceability can help find the convergence point; the common location that all of the contaminated food came through. Then authorities and the operator of that farm or facility can diagnose the problem further, identifying the exact source, such as an improperly sanitized shredding machine in a processing plant. Once the source of contamination has been identified and rectified, forward traceability can be used to identify where all the tainted items ended up. In practice however, much of the industry uses paper or manual systems, so the actual traceback process is very slow and fraught with error.[5] In fact, this inadequacy was a key impetus for including traceability in the 2010 Food Safety Modernization Act (FSMA).

Traceability Events and Data

The FSMA (Food Safety Modernization Act) directs the FDA to improve traceability, starting by running a set of traceability pilot projects. The final report from those pilot projects recommends that data be recorded for each of the following events as produce travels through the supply chain:

- Ship—sending produce to a customer
- Receive—receiving produce from a supplier
- Transform/process—e.g. clean, shred, bag
- Consume or dispose—sold to consumer or disposed

Transform/process events (e.g. preparing a bagged or boxed salad mix) will have two separate set of records: one for the inputs and another for the outputs. Data such as the items shipped/received, quantities, and lot numbers are to be recorded for each event (see sidebar at right). With this data, produce can be tracked at each step from origin to final consumption or disposal.

Recommended Event Data Elements

For each event, the FSMA pilots report recommends recording the following data:

- Event owner (firm submitting the info)
- Date, time, location of event
- Trading partner
- Item
- Lot/batch/serial #
- Quantity, unit of measure

And event-specific data, such as:

- PO #
- Work Order # (for processing events)
- BOL # (for shipping events)
- Carrier ID (shipping)
- Trailer ID (shipping)

[5] Manual 1-up/1-back systems are also open to abuse, since paper or privately held electronic records are easy to generate or falsify after the fact.

Industry Traceability Standards

The FSMA specifically does not say what technology to use and only mandates 1-up/1-back traceability. When contamination events occur, this leaves government agencies and agents with the work of contacting each party in the chain to request their 1-up or 1-back trace information, in order to stitch together an end-to-end picture. Companies have 24 hours to respond to such requests. The FSMA pilot report encourages the use of industry standards, which makes it easier to correlate the various records received from the different parties in the end-to-end chain. Traceability standards have been developed by GS1 working with PTI (Product Traceability Initiative). The standards focus on the item, quantity, lot/batch/serial #, harvest date (lot and harvest are optional). This data is put on a standardized label on each item, case, and/or pallet. Figure 1 below shows the 1-up/1-back approach using proprietary data and using industry standard data.

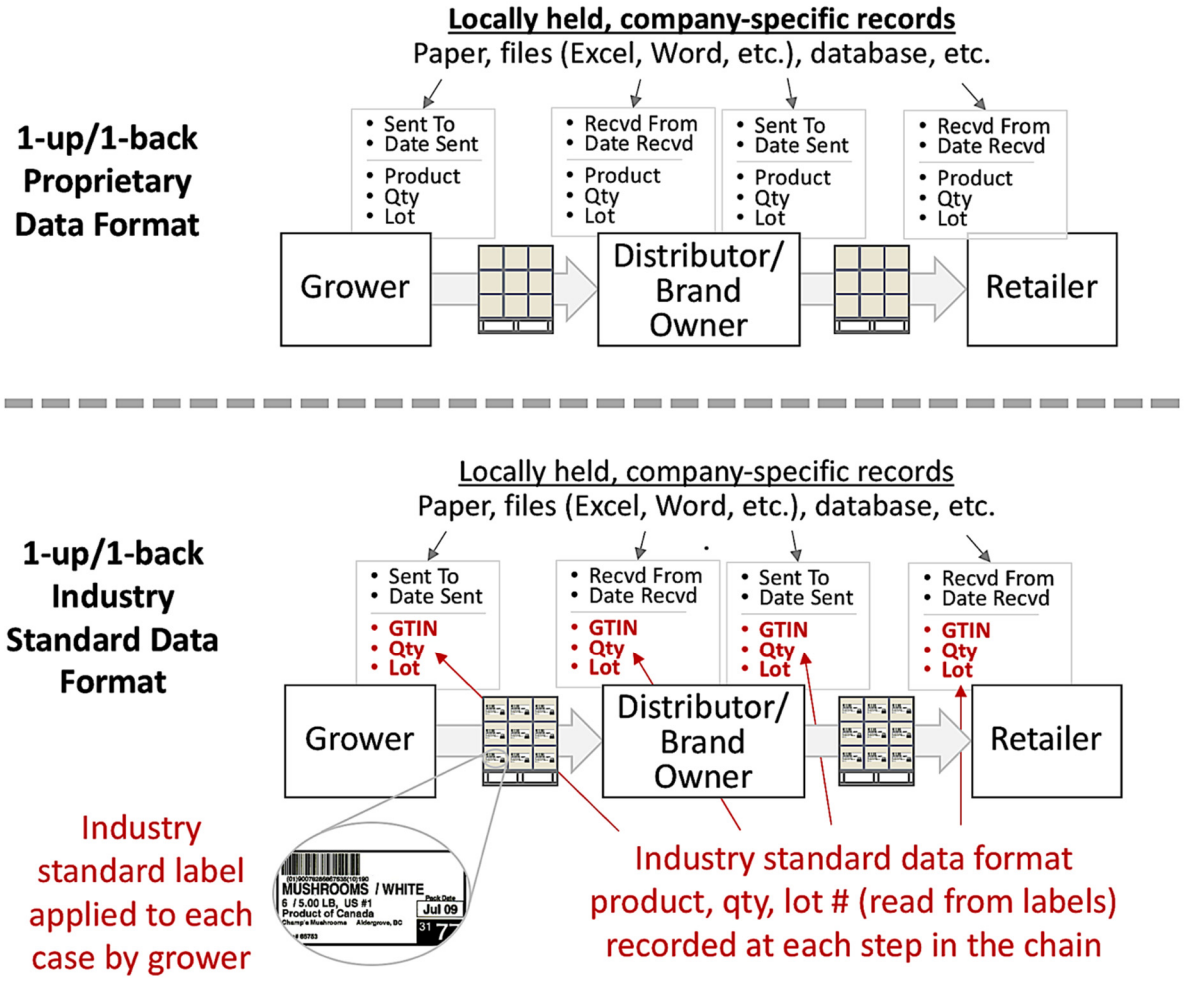

Figure 1 - 1-up/1-back Traceability—Proprietary vs. Standard Data Formats

The Role of Networked Cloud-based Systems and Blockchain

While not mandated by regulations, a networked cloud-based (aka SaaS) system provides a far better solution than 1-up/1-back. First, it is important to distinguish between Enterprise SaaS vs. Networked SaaS

architectures. Both have a shared code model (run in a single multi-tenant instance), but the enterprise SaaS model lacks a shared data model, where things like supply chain events can be written and shared across a network of trading partners.[6] In a networked SaaS system, the chain-of-custody data (hand-offs of batches or individual items from one party to the next) are recorded in a shared networked database, instantly accessible to authorized parties. This makes the trace-back and trace-forward processes virtually instantaneous. End-to-end tracing that used to take days or weeks can now be done in a couple of seconds. A blockchain architecture can bring similar network-wide data-sharing benefits.

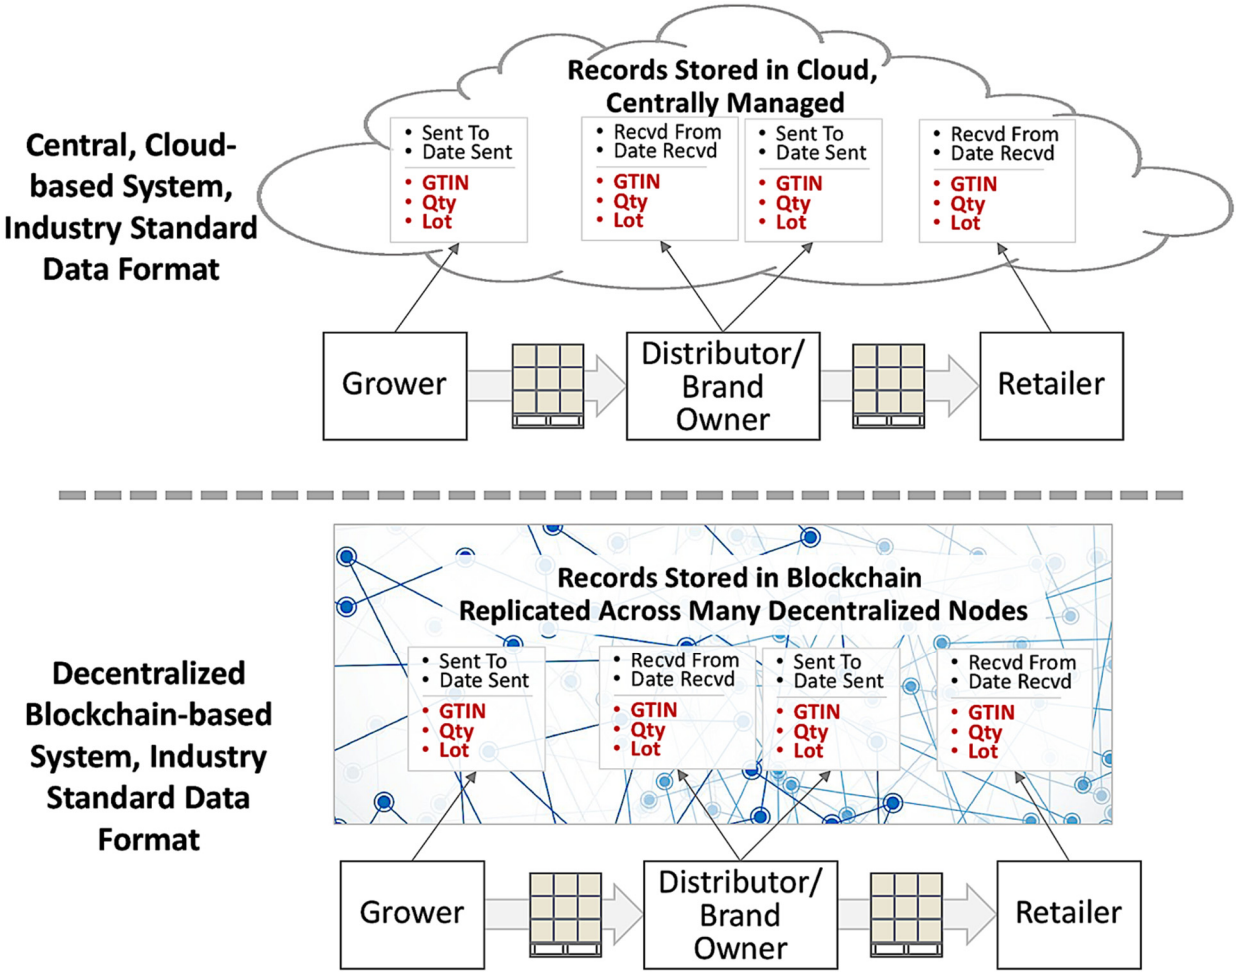

Figure 2 - Centralized SaaS/Cloud-based vs. Decentralized Blockchain-based Architectures

Figure 2 above illustrates the difference between a cloud-based vs. blockchain-based approach. In the cloud-based approach, a single central authority controls the system. The nodes may be physically distributed (and often are, to increase robustness and availability), but control is centralized. In contrast, control of the blockchain nodes is decentralized, allowing many parties to participate in ongoing verification and validation of the data. The blockchain architecture makes it much harder for someone to change data after the fact. We will further explore the differences between centralized cloud vs. blockchain below.

[6] For more details on the difference between Enterprise SaaS and Networked SaaS, see *Networked Platforms--Solutions for Multi-Party Inter-Enterprise Processes*.

Freshness and Safety

Data/Capabilities Required for Improving Freshness

Arguably the most valuable improvements for the produce supply chain come from increasing freshness and safety. Growers and retailers are always looking to reduce waste and spoilage in the supply chain and provide produce that has a longer post-purchase shelf life, with superior freshness. Improving freshness and reducing spoilage requires a number of additional data elements and capabilities, beyond those needed for traceability for provenance and recall:

- *End-to-end temperature monitoring*—Continuous tracking of the temperature of each pallet or case of product, starting at the point of harvest in the field all the way through delivery at the retail store. Without this information, accurate remaining freshness cannot be calculated.
- *Condition-based Expiration Date*—Based on the temperature exposure history of a given pallet of produce, a precise expiration date can be calculated that is much more accurate than the average shelf-life estimates based purely on days-since-harvest, that don't take into account the temperature exposure history of the pallet.
- *Temperature Response Profiles*—A model for predicting the Condition-based Expiration Date must be tailored to each produce variety, location, and time of season if it is to have any validity and accuracy. This requires testing and observing how each variety responds when exposed to different temperature histories. Without this, a model can generate only a crude estimate of the remaining shelf life, even if it has the end-to-end temperature history.
- *Per Customer/Location Requirements (Transit-time and Days-of-Freshness)*—In order to match each pallet to its optimal destination, the system needs to know the transit time and days-of-freshness requirements for each customer order.
- *Prescriptive Optimized Logistics (aka Intelligent Routing)*—Armed with the above pieces of information—the end-to-end temperature history, an accurate Condition-based Expiration Date, and per customer/location requirements—the system can tell line workers exactly what to do with each pallet or case at each key decision point in the supply chain. This includes local decisions, such as which pallets to put into the precool next, and intersite decisions, such as which pallets to send to which customers.

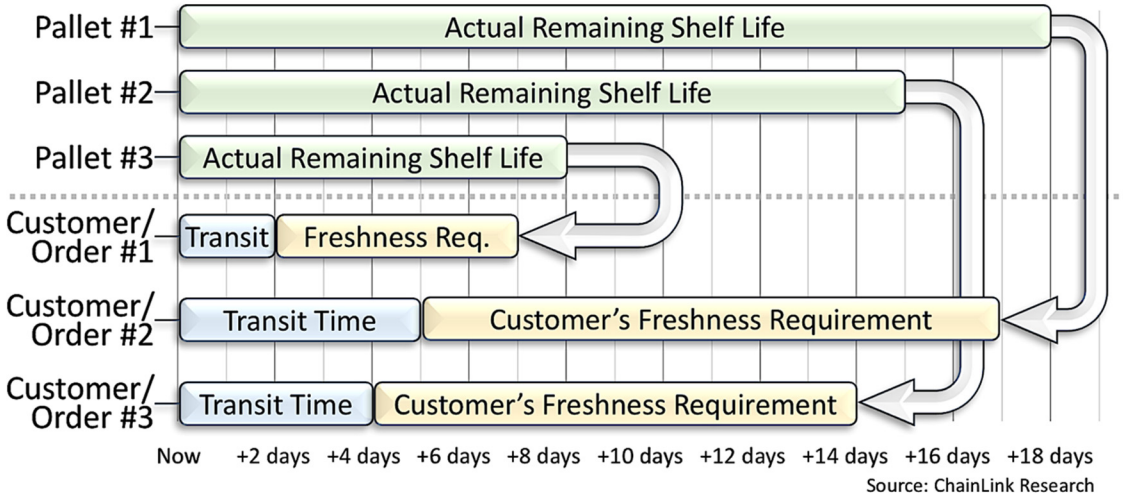

Figure 3 - Intelligent Routing—Matching Remaining Shelf Life to Customer Requirements in Distribution Decisions

For more details on these elements, see *Measuring Produce Freshness: Ensuring Delivered Freshness*. Without all of these data and capabilities, it is difficult for a system to make a material impact on improving freshness. The only company we're aware of who has brought all these pieces together is Zest Labs. They automatically collect pallet-level freshness data using wireless temperature tags (which they call *ZIPR tags*) that farm workers place into each pallet as soon as each pallet is loaded with produce out in the field. The tag starts recording temperature immediately and throughout the end-to-end journey. The system has algorithms to detect delays in turning on the temperature recording device, to ensure that any time spent sitting out in a hot field is included in the temperature record.

The Zest solution generates a Condition-based Expiration Date, which they call the ZIPR Code. This date is dynamically calculated (potentially changing as the pallet is exposed to different temperatures) based on extensive Temperature Response Profiling that Zest has conducted on many varieties of produce for different locations and time-of-harvest. They use AI-based and Machine Learning predictive and prescriptive analytics to dynamically determine freshness and enable prescriptive intelligent routing.

Ensuring Safety

While robust recall capabilities are important, they are reactive after-the-fact measures. Growers, brand owners, and retailers all want to ensure that the proper processes and procedures have been followed, to minimize (or ideally nearly eliminate) the chances of bacterial contamination. It is much better to proactively prevent adverse incidents from happening in the first place. Best practices are well understood and defined in HACCP[7] processes and standards. Skipping steps or taking shortcuts in these processes adds risk. The same system used to track provenance and freshness can and should be used to track HACCP test results and ensure that all process steps have been correctly executed.

A large grower or distributor operation may process thousands of pallets per day. During peak harvest season, these facilities are running at full capacity and then some. These are fast-paced environments, where things can go wrong. During peak season, any glitch puts tremendous pressure on the operation—whether it's a piece of machinery broken, a key employee or two out sick, extra hot weather requiring longer pre-cool times per pallet, or trucks arriving late. Growers and processors will tell you, off-the-record, "*When things go wrong and we get slammed, we will always find a way to get food through the system. And during peak season we get slammed every day.*" That is one reason it's so important to validate that HACCP requirements were met.

The Zest solution mentioned above outfits each pallet with a wireless ZIPR tag so that its movements and process steps can be automatically recorded, without requiring the workers to interrupt what they are doing. The system has built-in analytics to ensure that all the right steps have been done, in the right sequence and time, with algorithms to detect gaming of the system. For example, during harvest, pallets in the field are normally built and loaded onto a truck at approximately 10-15 minute intervals, depending on the type of produce and experience of the crew. If someone accidently or intentionally forgets to start recording temperature as soon as each pallet is built, the system will sense too little time between the pressing of the 'start recording' button for different pallets from the same field. It will generate an alert/warning to indicate that temperature recording likely was started too late for some or most of those pallets. Similarly, it can

[7] HACCP = Hazard Analysis and Critical Control Points, a systematic preventative approach to food safety.

detect if the temperature changes too rapidly, an indicator that someone is likely doing something suspect with the temperature sensor, such as putting it in cold water.

Similarly, the system checks that the right sequence has been followed, from field through yard acceptance, precool, and all of the steps. This way it can detect missing steps or meddling. It can also record HACCP processes, such as regular cleaning of the equipment, and results of tests such as lab analysis of the water used to clean any processing equipment that touches the product (such as the blades on a slicer). A positive test result indicating the presence of potentially harmful bacteria will automatically trigger an alert. Since the system has a complete record of where all product has gone, a retrieval can quarantine those pallets, stop them from being sent any further in the supply chain, and potentially pull them back.

With the Zest Fresh system, companies can use their own HACCP standards and processes, combined with the Zest solution, to proactively hold distributed pallets when an event has been detected, as well as record it all in a log accessible to the FDA upon request. Companies can implement their own policies in what to do when specific types of events happen, such as a HACCP step is missed or a test comes back positive. The policy can be implemented in a smart contract or off-chain (more on that below). All of this data (test results, events, etc.) can also be stored in a blockchain if stronger proof of non-tampering is required.

Blockchain's Role

Blockchain-specific Capabilities

Everything that has been described above can be and is done today by a centralized cloud-based system without the need for blockchain. However, recording the various transactions, HACCP steps, and temperature readings onto a blockchain can add trust and additional capabilities to the system:

- *Increased Security and Higher Availability*—With decentralization and embedded encryption, blockchain can be much harder to hack or bring down, provided there are a sufficiently large number of participating nodes executing the validation and storage. This strength comes at a cost, as there is now much more redundant storage and encryption processing required *(more on this below)*.
- *Tamper-resistance*—Blockchains are virtually tamper-proof once the data is written and validated. The confidence in the data thereby rests on the strength of the frontend consensus algorithms used and any protections in place to ensure valid data is being written in the first place. One approach is to have the data digitally signed by the edge devices originating them, and ensure those devices are fully secured[8] with secure boot and update mechanisms, trusted path/encrypted data streams, access controls, and anti-tampering mechanisms in place.[9] Another is algorithms that detect malfeasance (as discussed earlier for tampering with temperature tags). These *irregularity detection algorithms* are critical to revealing potentially inaccurate readings caused by intentional gaming of the system (like putting a temperature sensor into cold water) or simply sloppy execution (like forgetting to turn

[8] For example, Zest knows the location of every edge device it uses, including which subnet it should be on. If a piece of equipment, such as an IoT tag reader, gets unplugged, an alert is generated. Mobile equipment, such as devices used to take measurement in the field, use GPS to verify that the user is at the right location.
[9] For more details on securing edge devices, see *The IoT Security Imperative: Device Security Requirements*.

on the temperature recorder as soon as a pallet has been built in the field) at any point in the chain. Without irregularity detection algorithms, you may know with confidence that the data on the blockchain was indeed written by an IoT device, but will miss indicators of physical tampering with the device.

- *Automation via Smart Contracts*—Most blockchains have the ability to automatically execute smart contracts when certain conditions are met. While a centralized networked platform can also provide automated execution, smart contracts inherently provide transparency for all parties to see exactly what the rules of automation are. In contrast, the code that drives the automation provided by a centralized platform is typically hidden behind the scenes. Instead, the system normally exposes simple configuration rules to enable users to control the automation. Provided the UI is well designed, these configuration rules are likely to be more intuitive to understand by a non-programmer than the code executed by a smart contract.
- *Settlement*—Blockchains offer the possibility for extremely low-cost payments to be made between buyers and sellers of produce and/or related services (such as transportation). Until cryptocurrencies become widely accepted, a blockchain will likely require the ability for participants to pay and receive payment in fiat currency, with some sort of guarantee or limit on the effect of fluctuations in cryptocurrency exchange rates.
- *Soft Claims*—Soft claims can be a source of irritation between retailers and growers. Soft claims are often communicated and acknowledged informally with no requisite paper trail or system-of-record tracking them. Furthermore, retailers sometimes suspect the grower is not always shipping their best product and may be hiding a few bad pallets amongst the good ones. Conversely, growers may suspect the retailer is using some of their soft claims to cover up the retailer's own inaccurate forecast or poor inventory management, when they simply ordered too many and can't really use all of the pallets they ordered. With traditional methods, there is no easy way to methodically track soft claims, or to prove or disprove any of these suspicions. This makes disagreements and disputes all too common, thereby harming the relationship. With a system like Zest, the rejected pallet count is tracked and automatically communicated as a soft claim notification to the supplier. It can also be recorded on the blockchain, along with the objective measure of remaining freshness that both parties agree and rely on.
  By recording soft claims onto the blockchain, along with the ZIPR Code of any rejected pallets, it removes any ambiguity about how many pallets were rejected and the reason why, which can help soothe this sore spot in relationships.
- *Direct/Spot Market*—A distributed ledger can be used to create Uber-like markets, matching buyers with sellers in near real time, without the need for a middleman. This can be done not only for produce, but also for the trucking services required to transport the produce. These types of direct markets can be done without blockchain, and in fact have already been done for both produce and trucking, although still accounting for only a small percent of total transactions. In both cases, the bulk of produce and transportation has continued to be bought via direct contracts, because such relationships help ensure continuity of supply and quality for the buyer and more reliable demand for the seller. However, direct markets can be useful for spot demand, balancing out last minute demand/supply mismatches. Blockchain has inherent characteristics to support these direct markets, namely transparency (full audit trail), consensus (key parties have agreed and digitally signed each

transaction), immutability (people can't tamper with the records), smart contracts (automation of payment when certain events occur, such as delivery and passed inspection), and settlement capabilities (with very low or no payment fees). It becomes even more useful when combined with a system like Zest that provides full transparency into processes and conditions. Retailers tend to use spot markets only as a last resort, since they are essentially flying blind, with no relationship and very limited ability to establish trust in the supplier and the quality of the product they are delivering. Freshness and safety data, stored on a blockchain, with the anti-tampering algorithms described above, can help bridge that gap of trust for spot market participants. This should not only make buyers more comfortable, but enable sellers to get full market price for their produce.

- *Auditability and Restricted Transparency*—Most blockchains for commerce will be permissioned *(see below)*, so it is not the case that there is full public transparency, where everyone can see everything. In practice, a centralized cloud provider may choose to provide the same level of data transparency as a blockchain, but the blockchain ensures that nothing has been tampered with, thereby increasing auditability.

Architectural decisions about what data and functionality to put on the blockchain vs. on the networked SaaS platform will be shaped by the *much* higher cost of storage and execution on blockchains. This tradeoff is explored further in the section *On-chain Smart Contracts vs. Off-chain Automation* below.

Importance of a Blockchain Agnostic Architecture

There are many different blockchain and distributed ledger technology stacks emerging,[10] such as Hyperledger Fabric, Ethereum (various flavors), R3, Corda, and countless others. We are still in the early days of blockchain and it is too soon to tell who the winners and losers will be. That is especially true for the produce supply chain, where growers and retailers are just starting to learn about this technology and figure out what role it will play for them. They have not yet committed to a specific blockchain technology or implementation. It may turn out that the industry settles into two or three different camps over the long run.

Furthermore, even if the same technology is used, there could still be multiple different permissioned blockchains used by different parts of the industry. We could very well end up with a separate blockchain being used by each major retailer, for example. Therefore, it is prudent for growers and retailers to select a solution that is blockchain agnostic; one that is architected to integrate with any of the existing blockchain technologies as well as potential future yet-to-be-developed ones. The platform should also be capable of supporting multiple blockchains integrated onto the same platform, to deal with the likely scenario that not everyone in the industry is using the same chain.

Public vs. Permissioned Blockchains

Assurance of the various trading parties' identities and control over the privacy of transactions are required for most commercial uses of blockchains, including produce supply chains. For those reasons, a permissioned

[10] Some of the major technology companies, notably IBM, Microsoft, and Amazon/AWS are offering various development tools and environments that leverage these various blockchain technologies. They all seem to be hedging their bets as well, offering 'all of the above' options for different blockchain technologies.

blockchain approach is needed, rather than a public blockchain. The precise definitions of these terms are currently up for dispute,[11] but the table below shows the characteristics of public vs. permissioned blockchain as we are using the terms in this paper.

	PUBLIC	PERMISSIONED
PARTICIPATION	Anyone can participate	By invitation only, vetted either by a central authority, consensus, or other criteria
ACCESS CONTROL	Anyone can read and anyone can write (subject to validation)	Read and write access may be restricted to protect privacy of data
IDENTITY	Pseudonymous	Participants identified, preferably strongly
CONSENSUS	Typically requires a majority of validator nodes	May be done by a smaller set of nodes, such as stakeholders and/or knowledge-holders
CONTROL OF CODE	Anyone can make changes, but a majority of nodes decide which to keep	May be centralized or controlled by a consortium

Table 1 - Public vs. Permissioned Blockchain Characteristics

For most transactions in a produce supply chain, the parties cannot be anonymous or pseudonymous. There is a need for ongoing communications about orders, logistics, quality, acceptance test results and rejection, settlements, and other problems. Therefore, the identity of the transacting participants needs to be known and verified. Furthermore, the data about orders, prices, transactions, shipments, and so forth needs to be kept private to the parties involved. For competitive and security reasons, these data cannot be made available to the general public.

A scalable consensus algorithm is needed as well. Since participants have been identified and vetted, the heavyweight consensus approach of bitcoin or other public blockchains is not needed. Instead, digital signatures verifying acceptance of data by a few key stakeholder and knowledge-holder participants is sufficient. For example, a wireless tag reader may read the temperature history of a tag on a pallet, write the data into an off-chain database, and then create a record on the blockchain with a hash of the temperature data and pointer to the data. Ideally this should be done with full end-to-end security (as Zest does), with bi-directional authentication and strong encryption for all tag-to-reader and reader-to-cloud communications. A networked SaaS platform may then read that data, run machine learning-based algorithms to check for potential fraudulent data, and verify that everything appears to be legitimate, and then sign its verification of that data. Finally, the system may ask a worker at that location to visually verify[12] that the specific pallet in question is at that specific location, with a certain number of cases of a specific type of produce.

[11] There is currently a lack of broad consensus on precise definitions for these terms (including also the terms 'private' and 'consortium' blockchains). That is not surprising given the early stage of discovery and experimentation we are at in deploying various flavors of blockchain, beyond cryptocurrencies. We expect that lack of consensus to diminish over the next year or two, as people start to settle on specific definitions. We saw a similar pattern when the term 'cloud' first started becoming popular. At first there was a lot of handwaving and no universally accepted definition of what cloud meant. Within a couple of years, reasonably precise definitions coalesced and became widely accepted for specific categories of cloud services, such as IaaS, PaaS, and SaaS.

[12] In practice, this would likely be integrated into the existing receiving process, rather than creating a new extra step.

Thus, consensus may be met with just a small number of checks being made to validate the data being written on the blockchain, thereby taking minimal processing power and time, compared with the heavy requirements of say bitcoin's blockchain. In designing a consensus algorithm, the cost of validation needs to be weighed against the likelihood and consequences of someone tampering with the data. The system could also be designed to let the end users configure the type of consensus required (in essence, defining who needs to sign off on a transaction).

Automation and Smart Contracts

On-chain Smart Contracts vs. Off-chain Automation

A blockchain may contain smart contracts[13] that trigger and execute at key handoffs and decision points for each pallet or case of produce flowing throughout the end-to-end supply chain from farm to consumer. These can be used to automate key transactions and decisions. However, it is important to understand that storing data and running smart contracts on a blockchain is many orders of magnitude more expensive than using conventional computing resources.[14] Therefore, business applications will usually store most of the data and automation off-chain, using the blockchain only where it makes sense. Smart contracts will typically be used to encode mutually-agreed, high-level business rules and transactions, where multiple parties want full visibility and the ability to mutually validate the execution of transactions. Lower-level or 'behind-the-scenes' automation and algorithms will likely be executed off-chain. Thereby, the automation described below is likely to involve a hybrid of networked SaaS systems working in tandem with blockchain technology and smart contracts, rather than just blockchain-based smart contracts alone.

As we are in the early days of development and deployment of full scale blockchain applications in the supply chain, people are still experimenting and discovering what data and logic belongs on the blockchain vs. off-chain. The picture will become much clearer over time. In the meantime, below we describe hypothetical divisions of labor (i.e. execution) between blockchain-based smart contracts and off-chain automation logic, based on what might make sense today and in the near future.

[13] Smart contracts are logic on a blockchain that can execute automatically when triggered by specific events. The contracts may transfer value (in fiat currency or cryptocurrency) between trading partners.

[14] The cost of storing data and executing smart contracts on a public blockchain like bitcoin can be thousands of times more than what it costs to store and execute the same data and logic on a regular computer, largely due to the highly redundant nature of the execution and storage, but also to the amount of encryption logic being executed. A permissioned supply chain can reduce those costs by orders of magnitude, and eventually may get to within 10X or so of the cost of running on a single local machine.

Freshness-related Smart Contract and Off-chain Automation Logic

At the retailer, smart contracts can semi-automate the acceptance process, assuring that only produce with adequate remaining freshness and safety is accepted. For determining freshness, a smart contract needs access to the data and capabilities described in the section above on *Data/Capabilities Required for Improving Freshness*. In this case, the calculations of remaining days of freshness would be done off-chain by the networked platform, while the smart contract on the blockchain might contain the specific terms of acceptance, such as minimum days of freshness required. Such a smart contract could be triggered, for example, when a pallet is received by the retailer for acceptance or rejection (see sidebar at right).

Smart contracts do not eliminate the need for physical inspection for damage, as the temperature history does not reveal mechanical damage to the produce, nor discoloration[16] or other visual defects. Nevertheless, this type of semi-automation does accomplish several things:

- Allows QA people to focus on what they do best, inspecting for visible physical damage.
- Automates freshness assessment based on temperature history, thereby providing a much higher confidence in remaining freshness compared to visual inspection alone.
- Provides scalability and notification for acceptance/rejection processes.
- Reduces disputes, paperwork, and errors.
- Speeds up payments.

Example Smart Contract: Retailer's Acceptance

A smart contract between the retailer and grower/ distributor might include one or more freshness thresholds specifying exactly how many days of shelf life are needed for each type of produce by time-of-year (potentially by location[15] or other factors as well). For example, there may be an 'auto-rejection level' threshold; all pallets with shelf lives below that level are automatically rejected, and a higher 'full-acceptance level' threshold, above which all pallets are accepted (contingent on passing visual inspection).

Pallets with shelf life in-between the auto-reject and full-accept thresholds might be automatically accepted, but with the retailer paying only 85% of the full-shelf-life price. Or for those 'in-between pallets,' the contract may be set up to check inventory and forecast levels and if that particular item is in short supply, then it may accept those shorter shelf-life pallets. Alternatively, for 'in-between pallets' the contract may request the off-chain system to send all the relevant information (e.g. shelf life of pallets, current inventory levels, forecasted demand, etc.) to the responsible buyer or inventory manager at the retailer to make the decision. These are just illustrative examples: the actual business rules in the contract can be customized to meet the needs of both parties, just as the clauses in a traditional contract may be customized by both parties to serve their interests.

Further back in the chain, off-chain automation may be used to implement intelligent routing, and decide where to send each pallet of produce to match its shelf life with each destination's requirements. Intelligent routing is described briefly above in *Data/Capabilities Required for Improving Freshness* and in more detail in *Pallet-level Monitoring: Maximizing Delivered Shelf life in the End-to-End Fresh Food Supply Chain*.

[15] For example, one location may consume some types of produce faster than another, and thereby be willing and able to accept shorter-shelf-life produce, as reflected in the smart contract for that location. We don't necessarily expect this level of sophistication in early implementations, but over time we envision smart contracts that factor in location-specific criteria, current demand and inventory levels, or other relevant factors.
[16] In the future, other sensor data and analytics may be added to check for additional issues. For example, machine learning photo or video analytics could check for incorrect color or visual imperfections and sort the produce accordingly.

Safety-related Smart Contract Logic

Goals for safety in the produce supply include: 1) ensure the proper execution of safety and cleanliness processes and procedures, 2) detect any contamination and stop its distribution at the earliest possible point in time. For the first goal (ensuring proper execution of safety processes), it is best if the SaaS system is prescriptive for each process step, such as telling the workers at a packhouse which order to cool the pallets in, providing reminders and checklists to clean specific pieces of equipment, providing reminders and labels for samples to be taken and sent to the lab, and so forth. Monitoring can be accomplished via a combination of human inputs to the systems (such as checking off a checklist or filling out a form or taking pictures) and sensor data such as temperature history data to tell whether each pallet was properly cooled or video analytics that recognizes when someone is cleaning a machine. While most of this would be performed off-chain by the networked SaaS system, in some cases, the blockchain might be used to record evidence that safety procedures were followed.

In order to catch tainted produce as early as possible (the second goal above), we could envision the SaaS system doing an automated check of HACCP test results (e.g. microbial sample testing) as soon as the results are securely entered into the system by the labs. As well, at each step in the supply chain, the temperature history of the pallet could be checked to see if it has been exposed too long to elevated temperatures and thereby potentially be unsafe or unfit to be shipped or accepted. With the microbial sample results, a certain level of microbes may trigger a halt to production on the affected line and a 'hold' order for all associated produce in the supply chain. Then, when each pallet is about to be shipped or has just been received, an automated check is made on whether that pallet is on hold or clear to proceed. For a discussion of which portions of this should be executed on-chain vs. off-chain, see the sidebar *Example Division of Labor for Safety-Related Smart Contract*.

Example Division of Labor for Safety-Related Smart Contract

A smart contract that executes upon receipt of pallets at the retailer, and possibly at other key handoff points in the supply chain, could include a check of safety-related information, in addition to the freshness and provenance data. This safety data check could include confirmation that all related HACCP processes were completed correctly, and all post-processing tests were negative. Updating the safety data about each pallet on the blockchain is best managed by off-chain analytics that evaluate the securely entered lab HACCP test results and correlate or link those results to all of the pallets processed during the time-period associated with the sample being tested.

HAACP test results can take days to get back to the grower or processor, during which time all of the potentially affected pallets will have moved on, potentially through numerous transactions, far removed from the source. If there is an issue with a particular batch, the system needs to search and find all affected pallets. Blockchains are not set up for efficient searching, but the SaaS system's database is. So, in this case, the association of test results and pallets should be done off-chain. Once those linkages have been found, then the results of the test and status of all those pallets can be updated on the blockchain. Then the smart contracts can execute properly, with simple logic, based on that updated status.

The ability to broadcast a 'hold' for all pallets processed from a specific location on a set date is a necessary capability for food safety. A hybrid networked SaaS and blockchain solution elegantly bridges the process of correlating data off-chain, and then managing individual item status in the blockchain. Maintaining the current status in the blockchain enables independent smart contracts to automatically evaluate the item or pallet status.

Automated Recall

An even higher level of microbes might trigger an automatic recall. Automated recall is made possible when the end-to-end traceability information resides on a single shared network platform, including contact information for all parties in the supply chain. A message can be sent to hold/quarantine, return, or dispose of specific pallets. With true end-to-end tracing, the system knows exactly where each tainted pallet has been sent and currently resides and can thereby tell the responsible people at each location precisely how many affected pallets they have and give the precise unique ID # for each of those pallets. However, this type of automated recall capability is dependent on widespread adoption across the supply chain, which we discuss further below.

The ability to do instant recall has been the motivation behind a number of the publicized trials of blockchain for produce supply chain. Nevertheless, most of this automation can be accomplished much more economically by a networked SaaS platform, without the need for a blockchain or smart contracts.

Removing Tainted or Unsuitable Produce Early

By putting in these types of automated checks at various points in the supply chain, it prevents produce with inadequate shelf life or safety problems from moving further in the supply chain, thereby reducing waste and increasing customer satisfaction and safety. This is highly valuable. Visual inspection is standard practice, but produce appearance only reflects visible deterioration, not the actual remaining freshness or microbial contamination. There may be only very few days of remaining freshness and the product may still look fine. As a result, produce proceeds further in the supply chain but some of it will spoil before it is consumed.

It can spoil on the way to the retailer, at the retailer's DC, in transit to the store, at the store, or shortly after being bought by the consumer. All of these have bad consequences. At a minimum, supply chain resources are wasted and the retailer has less quantities of that variety of produce than they expected, potentially leading to shortages at the stores. In addition, it creates additional paperwork (for chargebacks), and potentially disputes between grower and retailer. Finally, a worst case is when the produce is actually sold to a consumer and spoils before its expected shelf life. You have an unhappy consumer and if that happens too often, the retailer (and potentially the produce brand owner) may lose that customer for life.

If the right intelligence is in place further back in the supply chain, then produce can be routed more intelligently, to locations that will actually consume it before the end of its shelf life. Waste across the system can be reduced, often by as much as 50%. This automation can all be accomplished by a Networked SaaS platform without blockchains or smart contracts. However, the addition of smart contracts provides a mechanism to automate acceptance and payment processes, in a way that makes the freshness criteria transparent to all parties involved.

Hybrid Systems Will Prevail

Blockchain technology alone cannot provide freshness, safety, provenance, and recall capabilities. That requires data and capabilities from outside the blockchain. It seems that the best emerging approach will be a hybrid consisting of 1) a centralized networked SaaS platform providing economical scalability and deep algorithmic and process capabilities, combined with 2) blockchain and smart contracts for transparency and validation. Blockchains are attractive because of their ability to create a shared, trusted single-version-of-the-truth between trading partners. However, a networked SaaS platform can provide a shared, trusted single-version-of-the-truth at a much lower cost.

Affordability is Key to Adoption

Affordability is essential for wide adoption and without wide adoption, end-to-end traceability cannot be achieved. A blockchain-based system can provide traceability, but only with full participation across the chain. Hence, the financial model (i.e. who pays) is important as it encourages or inhibits broad participation. In past attempts at traceability (such as PTI), the grower pays and yet gets very little actual benefit. A system that distributes costs and benefits more equitably across the chain is more likely to succeed.

Ensuring Validity of Data Being Written on the Blockchain

Furthermore, a networked SaaS platform provides a much higher level of confidence in the data written onto a blockchain. Take the case of a smart contract being executed when a pallet of produce is received at the retailer. If this is done with blockchain alone, then the contract might require a sign-off on the shipment by the retailer's QA person. If we also have a networked SaaS platform with end-to-end temperature tracking and produce temperature profiling capabilities, then the contract can include days-of-remaining freshness in the acceptance criteria. The SaaS platform will have confirmed that the temperature data read from the RFID tag has not been tampered with and is valid because all along in the journey of that pallet (and its tag) the platform has been doing process checks and irregularity detection, greatly reducing the chances there was any cheating or gaming of the system. As a result, the acceptance or rejection decision is made with much more confidence and intelligence about the actual condition of the produce.

Example of An Existing System Providing Hybrid Functionality

When looking at a hybrid networked-SaaS/blockchain solution, it is clear that the vast majority of the development effort and differentiated intellectual expertise is actually in the networked-SaaS platform. One solution provider, Zest Labs, offers virtually everything required for a hybrid solution, combining the best of networked SaaS and blockchain technologies. Working closely with some of the nation's largest grocers and growers, Zest has accumulated and created a depth of intellectual property over the past decade+ that includes:

- End-to-end temperature tracking, including the sensor and reader hardware, cloud-network, and, most critically, integration into all the various process steps from harvest in the field, to pre-cooling and other steps in the packhouse, transport, distribution, and receipt by the retailer;
- A rich set of temperature response profiles developed via extensive testing in their labs for different produce varieties, locations, and time-of-season combinations;

- Corresponding freshness models for all those permutations to calculate Condition-based Expiration Dates that accurately predict remaining shelf life;
- The complex processing mapping and distribution logic needed to do intelligent routing of pallets (sending each pallet to the destination best matched to its remaining shelf life);
- A system of control for growers to ensure that freshness and safety processes are being properly executed, resulting in longer shelf life pallets and fewer recalls;
- Very simple, unambiguous user interfaces refined over years to be useable by laborers and supervisors across the produce supply chain, in locations such as fast paced and chaotic packhouses;
- Deep analytic capabilities that let supervisors and managers make rapid optimal decisions when there are tradeoffs to be made or corrective action is needed;
- Irregularity detection analytics that understand and detect when processes are not being properly followed and/or there may be gaming of the sensor data;
- A deep understanding of industry dynamics and the needs of growers, retailers, and distributors—their ways of doing business, issues and concerns, typical processes and conventions, interests of each party, and so forth;
- Experience making all of this work out in the field (literally), overcoming all the various speed bumps and hurdles encountered in the real world.

No other solution provider we are aware of has pulled together a set of capabilities like this. The addition of blockchain technologies to Zest's solution is one more important capability to add to this list.

In the end, we predict that hybrid systems combining a centralized networked SaaS platform with blockchain capabilities will prevail. These will bring true transparency and trust. They are based on credible and accurate remaining freshness calculations, transparent monitoring of proper completion of all safety process steps, early detection of tainted produce requiring recall, the ability to prevent bad produce (either tainted or with inadequate shelf life) from proceeding further in the supply chain, and automated acceptance processes. When the costs and benefits for such a system are equitably distributed, we expect to see wider adoption and major benefits for everyone up and down the produce supply chain, all the way to the end consumer.

About ChainLink Research

ChainLink is a recognized leader in custom research and advisory services, with a focus on supply chain, Internet-of-Things, and blockchain. Founded in 2002, our emphasis from the start has been on inter-enterprise interactions and architectures ('the links in the chain'). We have conducted over 75 primary research projects in which we interviewed and surveyed over 5,000 professionals. Much of our research focuses on industry-specific use cases, business cases and ROI, and drivers/inhibitors of technology adoption and business change. As a result, we have developed a deep, multi-industry practice, founded on real-world, validated, supply chain-wide, end-to-end perspectives that have helped our clients understand, plan, and succeed as they move into the future.

321 Walnut Street, Suite 442

Newton, MA 02460-1927

617-762-4040

Email: info@CLResearch.com

Web: www.ChainLinkResearch.com