



Fog Use Case Scenarios

Use Case: Subsurface Imaging for
Oil & Gas Exploration

Vertical: Energy

An OpenFog Consortium Architectural Use Case

1 Snapshot: Fog-Enabled Oil & Gas Exploration



WHY FOG

Why is fog the best architecture for this use case?

Subsurface imaging in oil exploration is just one example of Energy use cases that require the deep distribution of intelligence that fog provides. Fog’s agility and programmability enable technicians in the field to adapt to conditions and to analyze results in real-time.



WHICH FOG PILLAR

Which fog pillar best describes this use case?

The fog pillar amplified by the Energy use case is Scalability. The imaging applications in particular require a great deal of local processing power and storage for terabyte data sets. Fog’s scalability enables real-time computation including support for complex compute algorithms.



VALUE

What are the business advantages of building this use case with fog?

The business advantage of fog-enabled oil exploration is that scientists are able to obtain much better data about underground geologic events and formations. In this use case, all computations and analysis are computed locally in the field without having to access the cloud or move media to a centralized laboratory. This use case describes a method for measuring the effectiveness of real-time activities such as drilling and fracking.



CLOUD & EDGE

How does this use case augment or supersede cloud and edge architectures?

Fog enables the bandwidth consumption for local processing in remote oil fields. The depth provided by fog hierarchies is key. Fog architectures provide a foundation for onsite power, data collection from thousands of sensors, and device to device communications. It also provides end-to-end security, preventing data theft.

2 Table of Contents

1	Snapshot: Fog-Enabled Oil & Gas Exploration	1
2	Table of Contents	2
3	Introduction	3
4	Fog Computing Overview	6
5	The OpenFog Reference Architecture	7
6	Benefits of Fog	8
7	Use Case Scenario: Subsurface Imaging	11
	Executive Summary	11
	Introduction	13
	Use Case	14
	Business Case	17
	Interoperability.....	18
	Advantages of Fog Computing	19
	Architecture	21
	Communication	24
	Testbed Considerations	25
8	Adherence to the OpenFog RA	27
9	Next Steps	28
10	About the OpenFog Consortium	29
11	Authors and Contributors List	30
12	Copyright / Disclaimer	31

3 Introduction

Note: The preamble section of this document (pages 3 through 11) is common across all OpenFog use cases. It provides descriptions and reference points for fog architectural attributes and properties. The Subsurface Imaging use case begins on page 11.

The [OpenFog Consortium](#) is defining applications and architectures for fog computing. The Consortium defines fog computing as: **A horizontal, system-level architecture that distributes computing, storage, control and networking functions closer to the users along a cloud-to-thing continuum.**

The first step in this architectural process is understanding the spectrum of vertical markets and applications that we expect fog computing technologies may serve. This document focuses on a representative use case that we believe spans many aspects of fog computing and therefore serves to define the functions we hope fog architecture, fog implementations, and fog deployments will provide.

It is important to understand how this use case fits into the overall process the Consortium uses to define interoperable and certifiable architectures. As shown in Figure 1, the use case described in detail in this document is a starting point for the suite of OpenFog technical documentation. When taken together, OpenFog use cases cover the basic fog functions of approximately 80% of the comprehensive set of IoT network applications we have identified for fog.

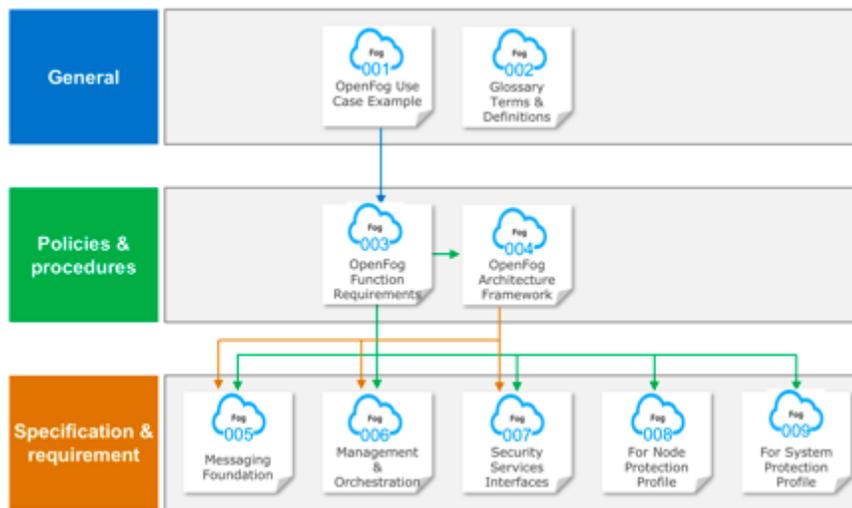


Figure 1. Hierarchy of OpenFog Consortium specification documentation

The composite of all use cases outlines a problem statement for OpenFog, describing the essential functions for all fog elements and networks. The Consortium extracts requirements from these use cases, and distills and correlates them to produce a detailed *Fog Requirements Document*. These requirements serve three important purposes:

1. To drive the OpenFog Reference Architecture;
2. To guide the development of OpenFog testbeds for testing and validation purposes; and
3. To provide guidance to implementers of fog nodes and networks.

The *Architecture Framework Document* is a compendium document that describes the key functional components of OpenFog as well as the interfaces between these components.

The Consortium also publishes additional documents, which describe details in areas such as security, management and orchestration, and messaging. Implementers may use the compendium as a guide for the conceptual planning and architecture design for their fog-based systems, and as implementation best practices for OpenFog elements and networks that will interoperate and can be certified as OpenFog compliant.

OpenFog Consortium workgroups reviewed and discussed hundreds of potential fog use cases spanning more than a dozen vertical markets related to IoT. The Consortium carefully selected a set of use cases that we believe spans a representative set of potential fog applications.

These use cases will highlight one or more representative attributes of fog such as latency, network bandwidth, reliability, security, programmability, scalability. The derived requirements from the use cases we include will cover an illustrative sample.

As mentioned, OpenFog technical requirements comprise a platform that covers approximately 80% of common fog functions. The remaining 20% of requirements needed to support specific use cases which are application dependent and won't be defined by the Consortium.

Readers should pay detailed attention to the subset of use cases that most closely match their areas of interest. We encourage you to browse additional use cases, as they may highlight less obvious aspects of fog that could prove valuable, and give insight into the rationale of the OpenFog requirements.

Readers are also encouraged to collect additional use cases and submit them to OpenFog for requirements extraction and potential inclusion in future use case documents.

4 Fog Computing Overview

Fog computing provides the missing link in the cloud-to-thing continuum. It is a critical architecture for today's connected world as it enables low latency, reliable operation, and removes the requirement for persistent cloud connectivity to address emerging use cases in Internet of Things (IoT), 5G, Artificial Intelligence (AI), Virtual Reality and Tactile Internet applications.

Fog architectures selectively move compute, storage, communication, control, and decision making closer to the network edge where data is being generated and used. This solves the limitations in current infrastructure to enable mission-critical, data-dense use cases.

Fog computing is an extension of the traditional cloud-based computing model where implementations of the architecture reside in multiple layers of a network's hierarchy. These extensions to the fog architecture may retain all the benefits of cloud computing, such as containerization, virtualization, orchestration, manageability, and efficiency.

The fog computing model provides the ability to move computation and storage from the cloud closer the edge, based on the needs of the data and the service requirements. These functions can potentially reside right next to the IoT sensors and actuators. The computational, networking, storage and acceleration elements of this new model are known as fog nodes. These nodes may also reside in the cloud, as they comprise a fluid system of connectivity and don't have to be fixed to the physical edge.

5 The OpenFog Reference Architecture

The OpenFog Consortium was founded on the principle that an open and interoperable fog computing architecture is necessary in today's increasingly connected world. Through an independently-run open membership ecosystem of industry, end users and universities, we can apply a broad coalition of knowledge to these technical and market challenges. We believe that proprietary or single vendor fog solutions are of limited value, as they can limit supplier diversity and ecosystems, resulting in a detrimental impact on market adoption, system efficiency, quality and innovation.

The [OpenFog Reference Architecture](#) (RA) is a medium- to high-level view of system architectures for fog nodes and networks. It is the result of a broad collaborative effort of the OpenFog ecosystem of industry, technology and university/research leaders. It was created to help business leaders, software developers, silicon architects and system designers create and maintain the hardware, software and system elements necessary for fog computing, as well as design, architect and develop solutions that enable fog-cloud, fog-thing and fog-fog interfaces.

6 Benefits of Fog

Fog computing targets cross-cutting concerns such as the control of performance, latency and efficiency, which are also key to the success of fog networks. Cloud and fog computing are on path to a mutually beneficial, inter-dependent continuum.

Certain functions are naturally more advantageous to carry out in fog nodes, while others are better suited to cloud. The traditional backend cloud will continue to remain an important part of computing systems as fog computing emerges. The segmentation of what tasks and single purpose functions go to fog and what goes to the backend cloud, are application and implementation/use case specific.

This segmentation can be planned and static, but can also change dynamically if the network state changes in areas such as processor loads, link bandwidths, storage capacities, fault events, security threats, energy availability, cost targets, and so on.

The OpenFog RA enables fog-cloud and fog-fog interfaces. OpenFog architectures offer several unique advantages over other approaches, which we term SCALE:

- **S**ecurity: Additional security to ensure safe, trusted transactions
- **C**ognition: Awareness of client-centric objectives to enable autonomy
- **A**gility: Rapid innovation and affordable scaling under a common infrastructure
- **L**atency: Real-time processing and cyber-physical system control
- **E**fficiency: Dynamic pooling of local unused resources from participating end-user devices

To illustrate this concept, let's look at a quick use case example: Consider an oil pipeline with pressure and flow sensors and control valves. One could transport all those sensor readings to the cloud (perhaps using expensive satellite links) to analyze the readings in

cloud servers to detect abnormal conditions, and send commands back to adjust the position of the valves.

There are several problems with this scenario: The bandwidth to transport the sensor and actuator data to and from the cloud could cost many thousands of dollars per month; those connections could be susceptible to hackers; it may take several hundred milliseconds to react to an abnormal sensor reading (during which time a major leak could spill significant oil); and if the connection to the cloud is down, or the cloud is overloaded, control is delayed or, in the worst case, completely lost.

Now, consider placing a hierarchy of local fog nodes near the pipeline. They can connect to sensors and actuators with inexpensive local networking facilities. These fog nodes immediately establish a community which provides the ability to collaborate. They can be highly secure, lessening the hacker threat. Fog nodes can also be given the authority to react to abnormal conditions in milliseconds, quickly closing valves to greatly reduce the severity of spills.

Local control in the fog nodes produces a more robust control system. Moving most of the decision-making functions of this control system to the fog – and only contacting the cloud occasionally to report status or receive commands – creates a superior control system.

Fog computing includes a set of high-level attributes of fog computing that we call the pillars; these include some of the fog advantages described in the pipeline control scenario. There are 8 pillars in total: security, scalability, openness, autonomy, reliability, agility, hierarchical organization and programmability. We will discuss all of these pillars in detail later in this document.

The OpenFog RA defines the required infrastructure to enable building Fog as a Service (FaaS) to address certain classes of business challenges. FaaS includes Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and many service constructs specific to fog. The infrastructure and architecture building blocks below illustrate how FaaS may be enabled; this will be expanded upon in the reference architecture document.

The OpenFog RA describes a generic fog platform that is designed to be applicable to any vertical market or application. This architecture is applicable across many different markets including, but not limited to, transportation, agriculture, smart cities, smart buildings, healthcare, hospitality, financial services, and more, providing business value for IoT, 5G and AI applications that require real-time decision making, low latency, improved security, privacy protection and are network-constrained.

7 Use Case Scenario: Subsurface Imaging

Use Case: Subsurface Imaging for Oil & Gas Exploration

Vertical: Energy

Executive Summary

Subsurface imaging and monitoring in real time is crucial for understanding subsurface structures and dynamics that may pose risks or opportunities for oil/gas and geothermal exploration and production. Subsurface imaging will also have applications in other industries, including government, aerospace, and even planetary exploration.

The integration of sensor networks with fog computing and geophysical imaging enables the creation of Real-time In-situ Subsurface Imaging (RISI) technology. Here in-situ means the data processing and imaging process are completed in the field instead of data being collected and then processes in the data center.

This technology can image 3D subsurface structures and dynamics in real-time, and search, identify and track subsurface targets. It has vast applications, such as:

- Surveying underground natural resources, such as oil and gas and geothermal
- Evaluating the security of national critical infrastructures and civil infrastructures
- Discovering subsurface hazards and targets such as landmines, hazardous materials and tunnels.

The subsurface imaging solution is built over an IoT network on which nodes must support cooperative sensing and computing. The IoT network is able to manage a large amount of geographically distributed nodes, which produce data that require different levels of real-time analytics and data aggregation.

Because of its distributed architecture and its real-time compute capabilities for massive amounts of data, fog computing is the only feasible enabling technology for real-time subsurface imaging and monitoring.

 <p>Challenges</p>	<ul style="list-style-type: none">• Large-scale and dense sensor networks add complexity and latency to management of graphics-based applications• Data-intensive computing involves significant real-time number crunching• Limited energy (e.g., battery life), bandwidth and computing resources at the site of data collection• Lack of real-time subsurface imaging to support timely decisionmaking• Unreliable connectivity due to remoteness and harsh environmental conditions• Expensive, maintenance-intensive process
 <p>Solution</p>	<ul style="list-style-type: none">• Fog nodes discover each other to self-form mesh networks• Fog nodes communicate with each other to compute the subsurface image• Fog nodes perform real-time streaming data processes• Distributed analytics algorithms running on fog are resilient in the face of failures of links or nodes.• Fog networks are self-adaptive to constant energy and bandwidth changes• Fog networks can be configured to transmit only important and final images to the cloud for more detailed analytics.
 <p>Technology</p>	<ul style="list-style-type: none">• Leverages distributed, cooperative fog computing and networking for subsurface imaging• Self-forming and self-organizing mesh networks• Situation-aware networking and computing

Introduction

The following describes an actual implementation of fog computing enabling real-time subsurface imaging. The purpose of presenting this use case is to promote more architectural conversations about fog computing use cases for the Energy sector.

For years, subsurface imaging has been used by energy companies for oil and gas exploration and production. Other sectors, such as government and aerospace, also use subsurface imaging.

The traditional subsurface imaging process involves massive data collection from hundreds and thousands of distributed seismic sensors. The data is collected and transferred to a data center for analysis and decision making. The collection process alone often requires manual retrieval, which adds to the tremendous time and cost of this process.

Deployment of sensors and network equipment in these inhospitable and often inaccessible environments adds to cost. Maintenance is also challenging: it can be difficult to know when and where problems occur in the collection chain. And then it can take time to restore worn or broken systems.

New wireless data acquisition systems can retrieve data from sensors to a local server for real-time processing onsite. But significant problems are still unsolved, including:

- Wireless data acquisition systems aren't designed for large-scale deployment; in fact, each sub-system is limited to small-scale deployments along a wire line.
- These systems require almost constant maintenance, particularly the need to replace batteries (often on a monthly basis).
- Because data is collected and processed in a central server, the infrastructure is expensive to set up.
- Security and privacy concerns related both to the unauthorized interception of data, and deliberate disruption of the network.

Utilities are looking for real-time subsurface imaging. First and foremost, real-time data collection and processing is the basis for faster decision making. This essentially means more processing at the point of data collection (in situ). But a real-time solution can also reduce the staggering costs of exploration and production.

Energy applications that depend on subsurface imaging directly benefit from the fog architecture.

Use Case

The Real-time In-situ Subsurface Imaging (RISI) network monitors and maps subsurface geophysical structures and dynamics. In situ means “in its original place” and refers to using computation, networking and storage resources of local fog nodes to analyze the sensor data without the need to transport it to the cloud.

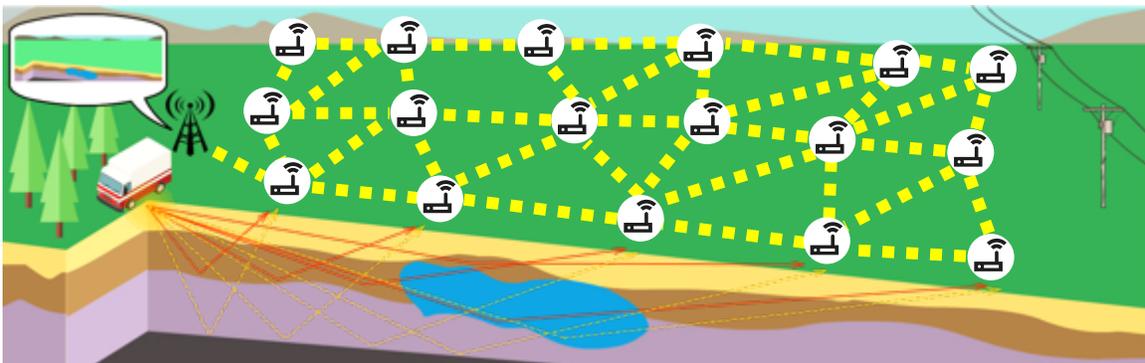


Figure 2. An illustration of Real-Time In-situ Subsurface Imaging (RISI). The wireless sensor devices are fog nodes that are placed on the ground and sense and process seismic signals from underground.

This wireless seismic network senses and processes seismic signals and computes 3D subsurface structures in situ and in real time. The system was developed by Sensorweb Research Laboratory and commercialized through Intelligent Dots.

The entire application infrastructure for RISI is enabled by fog computing. As shown in Figure 2, RISI consists of a mesh network of

sensor devices (fog nodes) that are placed on the ground and capable of sensing and processing seismic signals (emitted/reflected from the underground), performing distributed computing while communicating with other fog nodes to generate a real-time evolving subsurface image. Each fog node has an instrument and a small box housing a tiny yet powerful computational unit (e.g. Beaglebone Black, Raspberry Pi) and a low power radio communication device. Fog nodes are equipped with a GPS, seismometer, a battery and a solar panel.

Instead of collecting data in the field and transmitting it to the cloud for post-processing, the distributed seismic data are processed and inversion computing is performed in the in-situ network. The evolving 3D image (i.e., 4D imaging – including the time dimension) is computed and delivered in real time for visualization.

The fog architecture monitors and maps the subsurface geophysical structures and dynamics in real time. Instead of collecting data for post-processing, distributed seismic processing and imaging algorithms are performed in the in-situ network in real time. This generates the constantly evolving 3D seismic image for visualization.

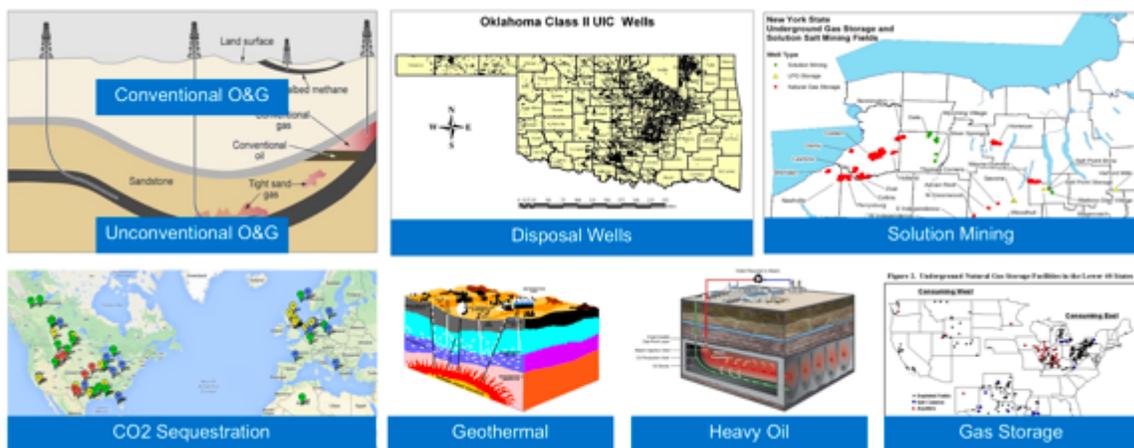


Figure 3. Fog computing enables more economical and reliable subsurface imaging in a wide variety of applications related to oil, gas and geothermal exploration and production.

RISI can detect, for example, the induced low-energy seismic events emitted from a reservoir during hydraulic fracturing stimulation. It provides real-time awareness (<1 second) of hydraulic fracturing

operations in unconventional oil and gas plays. This would be impossible without the continuous connectivity and processing power of the local fog.

Important decisions and analytics are executed inside fog nodes, instead of waiting until all information is sent to the cloud. This helps real-time monitoring systems to obtain timely data for decision making.

Intelligent fog computing and networking techniques distribute the data processing and inversion computing load among all the stations within the bandwidth and energy resource constraints of the network. The low latency of fog enables fast response in critical monitoring situations.

Fog nodes can be low-cost computing units (e.g. Raspberry Pi or BeagleBone Black) equipped with an internal, industry-standard processor. As shown in Figure 4, each fog node is equipped with a smart energy management module with its own solar panel and rechargeable battery for perpetual operation under normal weather conditions.

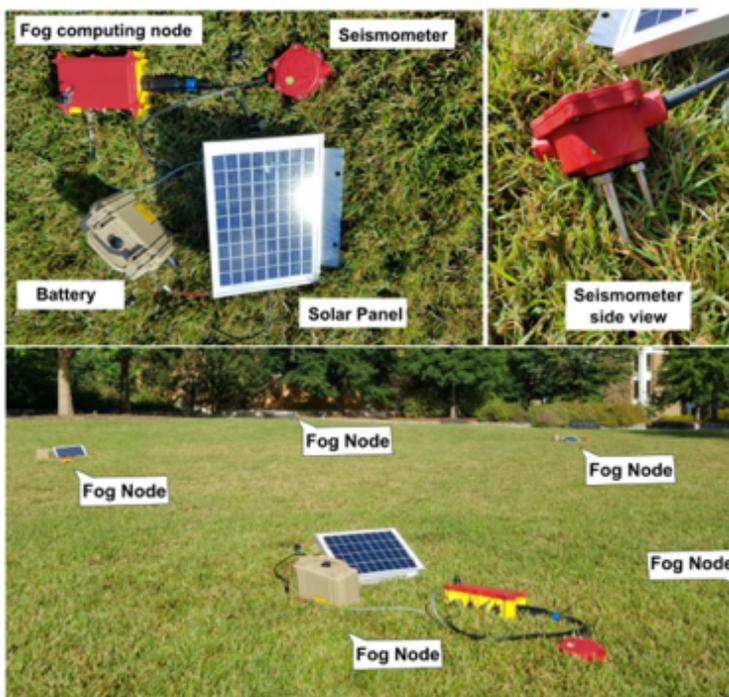


Figure 4. RISI fog computing node and network prototypes under deployment.

The actual sensors and the local fog nodes that process their data streams are basically like the heads of nails. The geophone is pounded a few inches to a few feet into the ground, but the node and its communications antenna and solar panels sit just at the surface. They receive, record and analyze patterns of vibrations, that can come either from a transmitter (the truck that “thumps” the ground to bounce waves off the layers of geology), or spontaneously arise in the structures themselves (micro-earthquakes).

To minimize the risk of data loss in remote environments where connectivity is intermittent, fog nodes utilize the limited communication, energy and storage resources collaboratively to maintain reliable operations. With fog, the failure of partial networks will not break down the whole system.

The fog network has built-in self-forming and self-healing capabilities. It requires no intervention during operation. When more nodes, sensors, or events are added, fog processing and computing will self-discover and self-adapt to compute higher resolution images with higher accuracy.

Business Case

Fog computing helps oil and gas companies understand how the reservoir responds to stimulation and its impact on customer economics. This reduces operations costs and mitigates environmental risks.

Consider monitoring a gas field that is being developed via fracking. To figure out the best place to drill, extensive seismic surveys are done, using the truck in the figure to thump the ground, and networks of hundreds or thousands of geophones to record the reflected vibrations, which are processed into 3D images. Geologists read the resulting pictures to figure out where to drill.

Then, as the drilling progresses, they monitor the images continuously to see if they are reaching the underground target areas. During the

fracking phase, the sounds collected by the network can help determine if they have pumped the right amount of fracking fluids in. Too little and the reservoir won't produce to its maximum (which can be costly). Too much fluid, and you have to pay more than you should for time, labor and materials or may contaminate groundwater – each of which is a costly proposition).

Fog computing lowers the cost and increases the efficiency and processing speed of the entire monitoring and mapping process, reducing compute-intensive processing from days (or even weeks or months) to seconds. This, in turn, leads to faster, more informed decision making (based on having up-to-date information).

Fog computing empowers fog nodes to process the raw data locally and periodically push the processed data to a central mainframe. This uses expensive bandwidth more efficiently.

Fog computing also reduces deployment and maintenance costs. Fog node deployment is a simple plug-and-play platform; it takes 10 seconds per station to deploy. Nodes don't require monthly battery replacement. A fog network is also extremely resilient, reducing the likelihood of failures in the field that require sending technicians into the field to identify and replace equipment.

The distributed nature of fog-based subsurface imaging accommodates a common hardware and software platform (for computing, networking, and storage), keeping pace with market growth and underscoring the scalability value of fog. For example, hundreds of inexpensive, mass produced fog nodes can be deployed quickly and seamlessly scale from small to big networks.

Interoperability

A fog node can also be an add-on to commercial data loggers, providing seamless integration with existing systems. With fog, there is functional interchangeability between fog nodes from different suppliers. This makes it easier for new equipment to interoperate with legacy equipment and brownfields. Additionally, it enables software

interoperability between many different producers with different domain expertise.

Advantages of Fog Computing

Fog computing greatly reduces latency and network bandwidth consumption. It greatly alleviates current issues with the availability of resources. The architecture for fog computing is based on eight pillars of elements, as identified by the OpenFog Consortium. In the subsurface imaging use case, the following pillars are addressed:

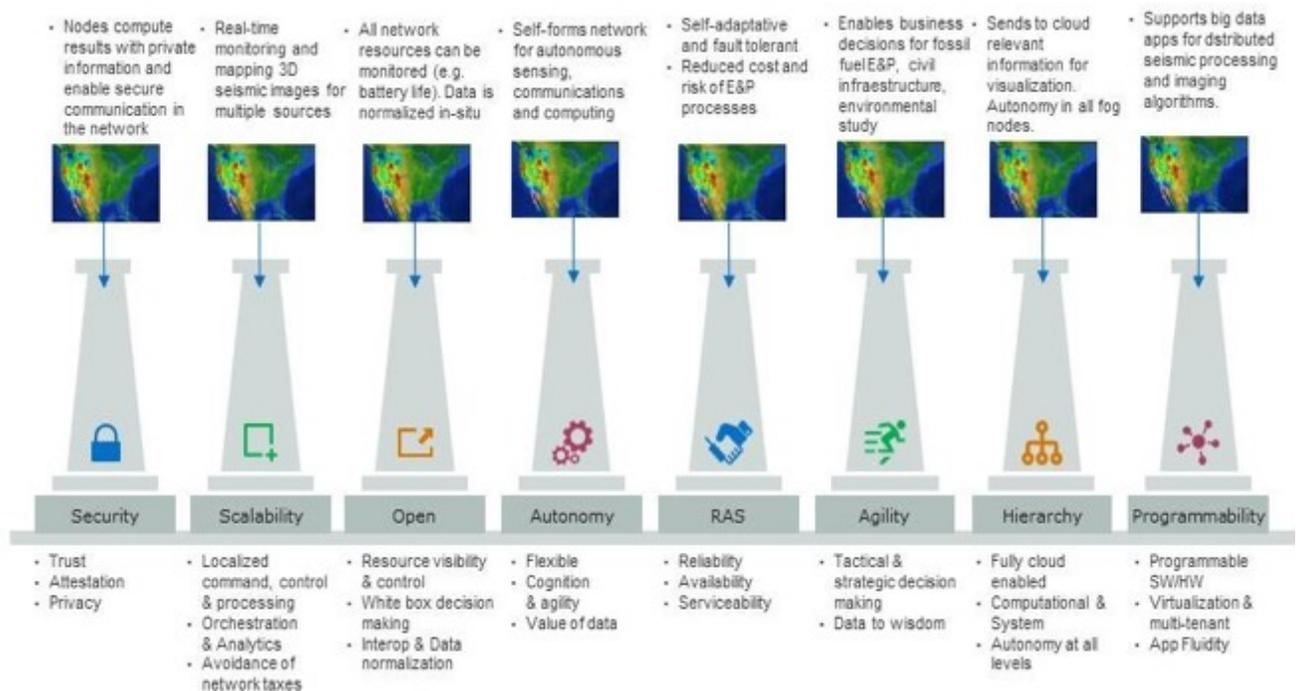


Figure 5. Subsurface imaging use case mapped to the 8 pillars of fog computing.

Security: Fog nodes securely exchange messages with each other, compute images in a cooperative fashion, and deliver the final subsurface image to the user. Fog also makes the network resistant to hacker attacks.

Scalability: The infrastructure of the fog nodes can be easily extended by adding more nodes, different algorithms, diverse communication types, and multiple outputs. This supports RISI’s ability to perform real-time monitoring and mapping of 3D seismic

images for multiple sources. The fog node will self-form a mesh network for autonomous sensing, communication and computing. More nodes can be added to form larger systems—or removed to form smaller systems—with little interruptions to current processes.

Open: The RISI framework grants open access to each fog node and monitors critical status (such as battery life, GPS signals, execution times, etc.). Additionally, the data is gathered, processed and normalized in situ. As the fog solution is open and interoperable, it should be possible to integrate equipment and software from multiple vendors and different generations into a seamless, unified network.

Autonomy: Fog provides RISI with the foundation of a self-forming network for autonomous sensing, communications and computing. Every node is self-ruling and can execute a wide variety of actions independently.

RAS: Every fog node is capable of accepting new interactions and tolerating failures in the network. This feature effectively can reduce costs and risks in E&P processes. This makes the RISI application self-adaptive and fault tolerant.

Agility: Due to real-time results, RISI improves the decision-making process in situ. The same network of fog nodes can be programmed and reprogrammed with different sets of applications and analytics algorithms depending upon the current exploration / drilling / or production phase of the project

Hierarchy: RISI transmits subsurface images to a server or cloud, enabling the user's decision-making. Every fog node knows its hierarchy in the network and performs external communication if needed. Some networks will use a hierarchy of fog nodes, with some fog nodes located at individual sensors, others associated with a well, and still others covering whole sections of a major field. Different sets of tasks are coordinated across these levels.

Programmability: The fog computing provides distributed cooperative computing. This enables RISI nodes to be remotely

programmable to support distributed data processing and computing algorithms.

Architecture

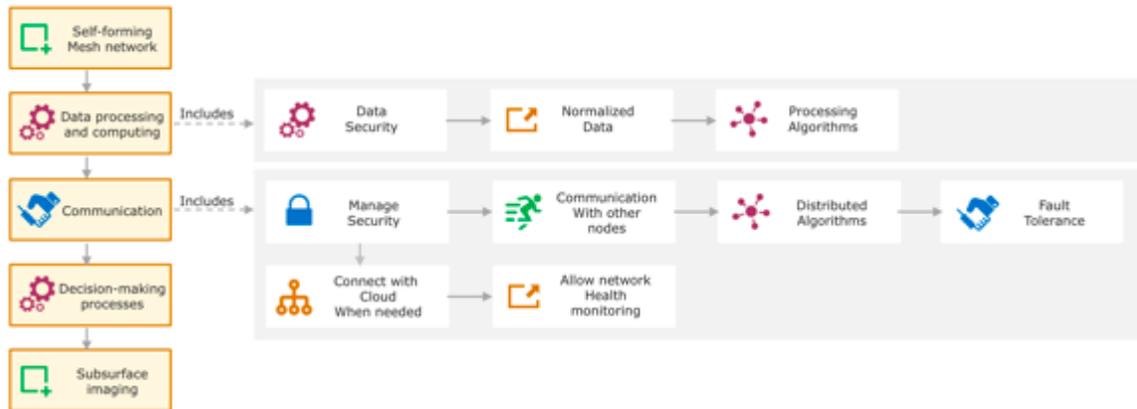


Figure 6. RISI technology and fog computing architecture.

The RISI system architectural design is shown in Figure 6. There are five main processes in the architecture to resolve the subsurface structure (self-forming mesh, data processing, communications, decision-making processes and final result generation). However, some main processes involve a series of non-trivial subprocesses that are needed to ensure a reliable, secure and timely system.

Self-forming mesh network. The deployment of RISI technology includes a self-forming mesh network, which enables scalability of the system. Adding new nodes or removing existing nodes can be accomplished without interruption.

Data processing and computing. The key aspect of the fog computing model for subsurface imaging is that it performs computation among the nodes of the network without relying on a central coordinator. Each fog node is able to collect, process and communicate important and filtered data with other fog nodes. All fog nodes work collaboratively to monitor and map the subsurface geophysical structures and dynamics in real time.

An example is the data correlation function between fog nodes. The data correlation is a crucial step in subsurface imaging because it helps to construct the subsurface impulse response between two sensors on the surface. Fog nodes correlate preprocessed data in real time under certain bandwidth constraints.

The data processing inside RISI technology includes a series of steps:

- **Data acquisition:** Every node acquires a real-time data stream autonomously from its attached sensor.
- **Normalized data:** Each fog node normalizes and prepares data for analytics.
- **Processing algorithms:** A key feature of RISI is that it provides the programmability to include and modify different types of algorithms for in-situ processing.

Communication. RISI technology enables different types of communication (e.g. wireless, X-bee). The communication process includes:

- **Security management:** Security is a critical feature of RISI technology. Communications between nodes and/or cloud are made using a secure channel to ensure data integrity, privacy and hacking resistance.
- **Communication with other nodes:** This feature enables the agility of the RISI technology. Communication with other nodes is faster and is performed inside the mesh network using communication range and bandwidth restrictions.
- **Distributed algorithms:** As described above, RISI technology provides programmability to include a variety of distributed algorithms that perform critical functions inside the system.
- **Fault tolerance:** With fog computing, RISI has built-in self-healing capabilities. The failure of partial networks will not break down the whole system, which provides the RAS (Reliability, Availability and Serviceability).
- **Connect with cloud when needed:** RISI technology allows fog nodes to connect with servers or clouds for off-line decision-making process, enabling a hierarchical vision of the system.

Without fog computing, processing subsurface imaging would be dependent on sending all data to the cloud, which means high latency, high bandwidth demand, less reliability and less fault tolerance.

- Network health monitoring: Even though the fog networks run autonomously and independently, RISI allows health monitoring of all nodes in real-time (e.g. battery life, GPS signal, etc.).

Decision-making processes. In RISI technology, every node is able to perform real-time decisions. This feature consolidates the vision of the autonomy of the system.

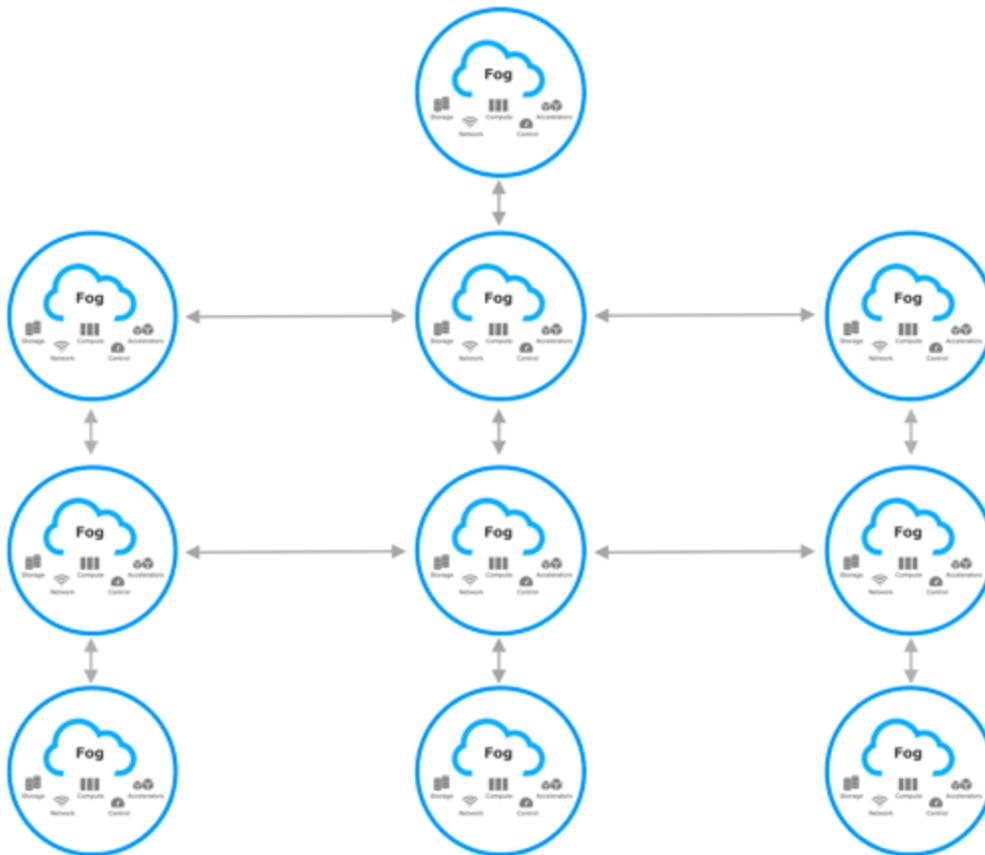


Figure 7. This illustrates a mesh network of fog nodes that process and cooperate in the computation of distributed algorithms for subsurface imaging. At the horizontal level, the system allows all fog nodes to cooperate on distributed and efficient computation. At the vertical level, some fog nodes send relevant data to the cloud for analytics that enable decision making.

Communication

Fog nodes generally are of most value in scenarios where data from thousands or even millions of devices must be collected, analyzed and acted upon in micro and milliseconds. In the subsurface imaging scenario, data analysis needs to be executed among fog nodes through message exchanges. Communication plays an important role in these kinds of applications.

Some key fog communication considerations for RISI include:

- Mesh networking protocols
- Distributed MAC protocols
- Unicast, convergecast and broadcast communications.

As shown in Figure 8, each fog node is equipped with a geophone to sense the subsurface activities. The data is then preprocessed by the fog node and converted in a suitable packet when the communication is needed.

The system is self-adaptive over communication, which means it chooses the protocol available in the network (e.g., UDP, TCP, or proprietary protocols). Results are obtained by using only node-to-node fog communication.

At any time, any fog node can send a summarized data product to the cloud for high-level decision making. The orange square in Figure 7 represents the life of a packet inside a specific fog node. However, all nodes perform the same strategy to obtain a subsurface image.

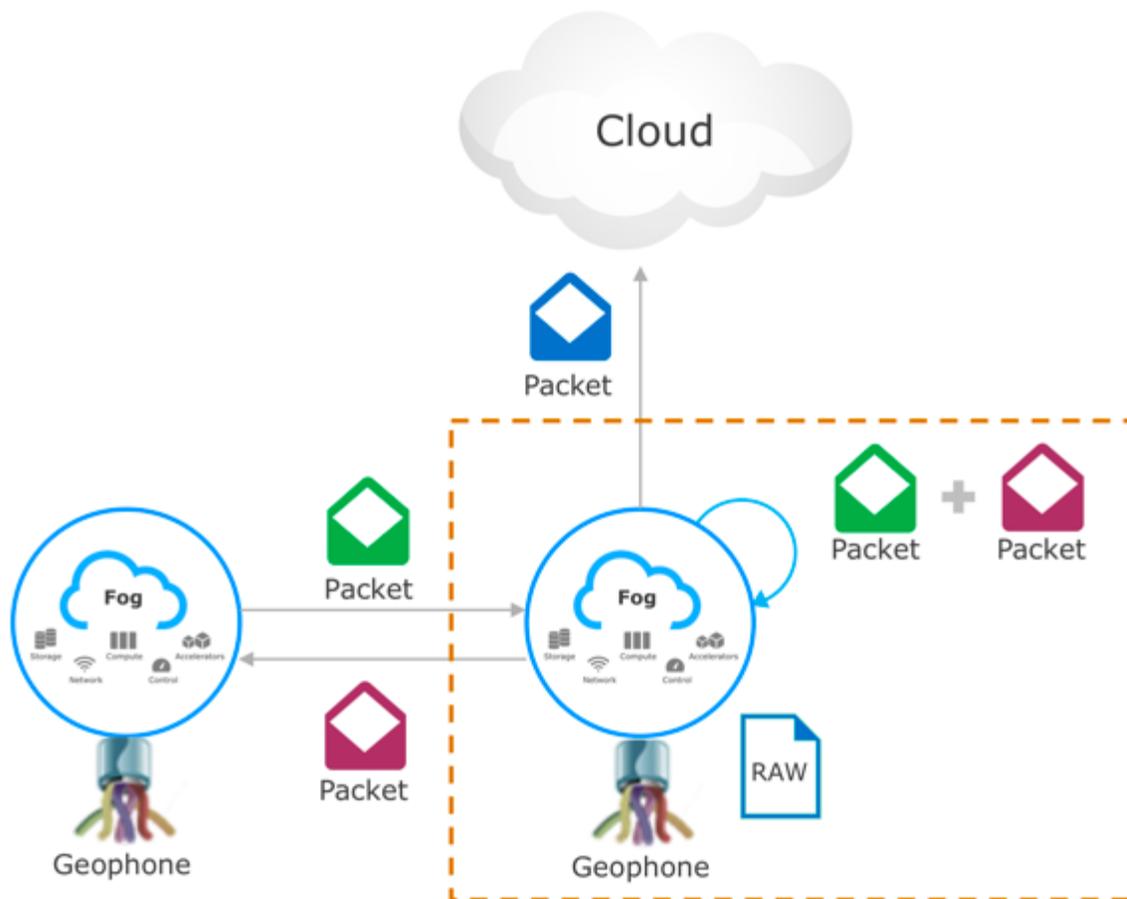


Figure 8. This illustrates the life of a packet during the communication process inside the system. Note that a node can act as a gateway to reach the cloud.

Some important considerations:

- The network should provide the scalability, availability, and flexibility required by the communication pattern or process
- The network should also provide whatever QoS is required to prioritize critical or latency-sensitive data and even guarantee delivery
- QoS should be addressed at fog node level.

Testbed Considerations

The hierarchy of OpenFog testbeds will be structured as follows:

1. Many small, research-oriented locations that OpenFog Members are able to access will focus on proving the high-level OpenFog architectural requirements and satisfying the minimum interoperability requirements via their Proof-Of-Technology(POT) Testbeds. The outcome of these Proof-Of-Technology testbeds could be open source code or a research publication available to OpenFog members.
2. Medium-sized, Interoperability Operation Model (IOM) testbeds will focus on overall solutions and end-to-end applications, with at least three OpenFog Sponsors participating to promote usage of diverse OpenFog Ready Solutions. They will demonstrate adherence to the OpenFog Reference Architecture and component-level interoperability and compatibility.
3. Large, regional testbeds will test pre-productization devices for application to the co-located OpenFog Certification Lab. After the OpenFog Certification Lab validates a product, members will be able to release it as an OpenFog Certified product. We expect many verticals, use cases, and individual applications will have specific requirements for interoperability and preferences for certain types of testbeds, and the Consortium intends to adapt to their needs.

8 Adherence to the OpenFog RA

The OpenFog Consortium intends to partner with standards development organizations and provide detailed requirements to facilitate a deeper level of interoperability. This will take time, as establishing new standards is a lengthy process. Prior to finalization of these detailed standards, the Consortium is laying the groundwork for component level interoperability and certification. Testbeds will prove the validity of the [OpenFog Reference Architecture](#) (RA) through adherence to the architectural principles.

9 Next Steps

The [OpenFog Reference Architecture](#) (RA) is the first step in creating industry standards for fog computing. It represents an industry commitment toward cooperative, open and interoperable fog systems to accelerate advanced deployments in smart cities, smart energy, smart transportation, smart healthcare, smart manufacturing and more. Its eight pillars imply requirements to every part of the fog supply chain: component manufacturers, system vendors, software providers, application developers.

Looking forward, the OpenFog Consortium will publish additional details and guidance on this architecture, specify APIs for key interfaces, and work with standards organizations such as IEEE on recommended standards. The OpenFog technical community is working on a suite of follow-on specifications, testbeds which prove the architecture, lists of requirements, and new use cases to enable component-level interoperability. Eventually, this work will lead to certification of interoperable elements and systems, based on compliance to the OpenFog RA.

For more information, please contact info@openfogconsortium.org.

10 About the OpenFog Consortium

The OpenFog Consortium was founded to accelerate the adoption of fog computing and address bandwidth, latency and communications challenges associated with IoT, 5G and AI applications. Committed to creating open technologies, its mission is to create and validate a framework for secure and efficient information processing between clouds, endpoints, and services. OpenFog was founded in November 2015 and today represents the leading researchers and innovators in fog computing.

For more information, visit <http://www.openfogconsortium.org/>;
Twitter [@openfog](https://twitter.com/openfog); and LinkedIn [/company/openfog-consortium](https://www.linkedin.com/company/openfog-consortium).



11 Authors and Contributors List

Authors	Contributors
WenZhan Song, University of Georgia	Chuck Byers, Cisco Systems
Maria Valero, University of Georgia	Judith Kelley, OpenFog Consortium
	Evan Birkhead, OpenFog Consortium

Note: All publicly available use cases are reviewed and approved by the OpenFog Technical Committee.



12 Copyright / Disclaimer

This reference document is designed to provide a foundation for extracting requirements when developing fog-based architectures. It is a compendium document to the OpenFog Reference Architecture. <https://www.openfogconsortium.org/ra/>

Copyright © OpenFog Consortium, 2017.