



# Fog Use Case Scenarios

**Use Case:** Smart Buildings  
**Vertical:** Smart Spaces

An OpenFog Consortium Architectural Use Case

## 1 Snapshot: Fog-Enabled Smart Buildings



### WHY FOG

Why is fog the best architecture for this use case?

The Smart Buildings use case features localized nodes for real-time computation at the room level, floor level and building level that are collecting data from sensors and operating actuators throughout the building. Without fog's inherent distribution, nodes on each layer would quickly be overloaded.



### WHICH FOG PILLAR

Which fog pillar best describes this use case?

Hierarchy is the OpenFog pillar most amplified by the Smart Building use case. The use case demonstrates how fog nodes at the room level, floor level, building level and cloud level can be hierarchically architected for efficient real-time processing, enabling dozens of new applications.



### VALUE

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

The business advantage of this use case is that its fog-driven attributes help optimize the security, efficiency and occupant comfort of the building, making it more cost-effective to operate and more attractive to tenants. Fog-enabled interoperability ensures that all infrastructure elements will work together, making it easy for operations managers to choose the best providers – from the smallest sensor up to the most complex cloud analytics application.



### CLOUD & EDGE

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

Fog-architected Smart Buildings provide a better security and privacy platform. Local processing also accounts for latency considerations for real-time applications such as credential scanners. Fog excels at bandwidth-intensive applications such as video surveillance, which would create bottlenecks in buildings with hundreds or thousands of cameras. Edge computing is not designed to manage the hierarchy of buildings, floors and rooms.

## 2 Table of Contents

<b>1</b>	<b>Snapshot: Fog-Enabled Smart Buildings</b> .....	<b>1</b>
<b>2</b>	<b>Table of Contents</b> .....	<b>2</b>
<b>3</b>	<b>Introduction</b> .....	<b>3</b>
<b>4</b>	<b>Fog Computing Overview</b> .....	<b>6</b>
<b>5</b>	<b>The OpenFog Reference Architecture</b> .....	<b>7</b>
<b>6</b>	<b>Benefits of Fog</b> .....	<b>8</b>
<b>7</b>	<b>Use Case Scenario: Smart Buildings</b> .....	<b>11</b>
	Executive Summary .....	11
	Introduction .....	13
	Flagship Use Cases .....	15
	Emergency / Disaster Response .....	16
	Optimizing an Employee’s Performance.....	18
	Advanced Services.....	21
	Advantages of the Fog Computing Approach .....	23
	Mapping to OpenFog Communications Architecture.....	31
	Smart Buildings in Smart Cities .....	33
	Fast Track to OpenFog Testbeds.....	34
<b>8</b>	<b>Adherence to the OpenFog Reference Architecture</b> .....	<b>36</b>
<b>9</b>	<b>Next Steps</b> .....	<b>37</b>
<b>10</b>	<b>About the OpenFog Consortium</b> .....	<b>38</b>
<b>11</b>	<b>Authors and Contributors List</b> .....	<b>39</b>
<b>12</b>	<b>Copyright / Disclaimer</b> .....	<b>40</b>

### 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 Smart Building 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:

- **Security:** Additional security to ensure safe, trusted transactions
- **Cognition:** Awareness of client-centric objectives to enable autonomy
- **Agility:** Rapid innovation and affordable scaling under a common infrastructure
- **Latency:** Real-time processing and cyber-physical system control
- **Efficiency:** 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: Smart Buildings

Use Case: Smart Buildings

Vertical: Smart Spaces

### Executive Summary

Today's smart buildings are beginning to leverage the Industrial Internet for improved business outcomes, such as better energy efficiency, improved occupant experience, and lower operational costs. They may contain thousands of sensors measuring various building operating parameters such as temperature, humidity, occupancy, energy usage, keycard readers, parking space occupancy, fire, smoke, flood, security, elevators, and air quality.

These sensors collectively capture massive amounts of data that must be transmitted, stored, analyzed and acted upon, often in real-time, to provide a truly smart building experience. These actions require thousands of actuators capable of exercising fine-granularity control over lighting, environment, security, safety, and building systems.

Some of this processing and actuating is extremely time-sensitive, and requires a real-time response from devices in close proximity to the edge. For example, turning on fire suppression systems in response to detecting a fire event and guiding occupants to the nearest exit, or locking down an area if an unauthorized person tries to gain entry.

Some building systems are life critical, and require higher availability than cloud-based solutions can achieve. Some smart building applications are so bandwidth intensive that they would swamp the building's fiber access bandwidth. Other solutions require huge amounts of historical data to feed computationally intensive machine learning models, and therefore can only take place in the cloud.

The fog computing architecture gives smart building technology suppliers the flexibility to collaborate with their customers to create more targeted, outcome-based solutions. By moving computation, networking and storage to locations within the building, it removes the constraint of operating entirely at the edge with no long-term learning abilities, or entirely in the cloud with inadequate real-time responses. This will help address the high volume of untapped opportunities in the market.



### Challenges

- Each building has a large number of interdependent systems governed by silo'd applications that do not share data.
- Rapid decisions about security, temperature settings, emergency responses, and more must be made on location in real time; latency delays ruin the "smart" experience.
- Smart buildings generate terabytes of data per day, which is unmanageable by the cloud alone, and can overload Internet access.
- Critical building systems require highly-reliable computation and bandwidth, which can be challenged in emergency situations.
- Buildings are vulnerable to security breaches and hacking.



### Solution

- Fog computing uses a distributed computing approach to create smart, connected spaces.
- Fog's ability to support compute-intensive applications locally enables real-time decision-making.
- Fog-based deployments provide the opportunity to build real-time, latency-sensitive operations paired with long-term, constantly improving experiences.
- Fog-based deployments break down silos between applications, systems and networks, enabling new experiences based on intelligent data-sharing and network-wide interoperability.



### Technology

- Fog nodes, which are the critical compute components in fog networks, intelligently partition data processing into layers of fog nodes between sensors and actuators on the edge and the cloud.
- Fog nodes for individual rooms can perform all monitoring and response functions. Fog nodes at the floor / section / wing level can coordinate rooms.
- Fog can lower infrastructure and equipment costs by combining functions from multiple overlay building networks onto a single fog network.
- High priority algorithms can be located on fog nodes close to the edge, for rapid latency and reaction to safety situations.
- Fog can extend the advantages of highly secure networking to large numbers of inexpensive smart building endpoints.

## Introduction

Buildings across the world are becoming smarter every day. The rapidly decreasing cost of computing, the growing demand for more energy, carbon and water efficient structures, and the increasing desire for omnipresent connectivity are all driving builders to build smarter. This means thousands of sensors per building, measuring everything from space temperature to people counting, generating terabytes of data per day.

Occupants don't find buildings smart if they just generate this data, however. The intelligence comes when buildings learn to act in real-time to delight their occupants, to customize experiences, to conserve their own energy use, and to manage the security and integrity of their data with the appropriate discretion. For these new capabilities to truly deliver, thousands of sensors and actuators must be managed, controlled, archived, and synchronized to incredible precision.

Dealing with these increasing demands on data generated by buildings requires a smarter architecture. A Smart Building built on today's model of proprietary edge devices managed by their vendor from the cloud would be incapable of reaching this kind of intelligence:

- Sending terabytes of sensor data to the cloud every hour would consume all of the internet bandwidth for the building and restrict high-data rate services such as virtual or augmented reality.
- Waiting for that data to be processed in the cloud before instructions are sent back to the actuators would imperil critical life-safety systems and dissatisfy users expecting seamless, dynamic operations.
- Managing a heterogeneous device ecosystem with no built-in interoperability, no local ownership, and separate, overlaid networks for HVAC, lighting, telecom, security, fire safety, etc. would be a nightmare for owners just trying to keep their building running – not to mention those trying to enable incredibly Smart experiences.

The OpenFog architecture is designed to solve these problems.

A Smart Building built on the OpenFog architecture will have enough local storage and compute that neither its emergency response system nor its augmented reality wayfinding system need to depend on a cloud internet connection to function. All of its devices, and the fog nodes governing them, will conform to open standards that allow for simple set-up, integration and management.

The Smart Building owner will actually be in control of what's in his or her building, and be empowered to deliver next-generation experiences to their occupants. The following scenarios seek to illustrate exactly what type of experiences can be made possible, and how the OpenFog architecture is uniquely designed to make them succeed.

### Flagship Use Cases

The fog-based Smart Building has many different autonomous systems working together to ensure its occupants are as comfortable and productive as possible. An examination of three specific use cases – 1) Emergency/Disaster Response, 2) Optimizing an Occupant's Performance, and 3) Advanced Services Dependent Upon OpenFog - demonstrates what will be required of the OpenFog architecture in order to enable this type of systemic cooperation.




Figure 2. A Smart Building designed with the OpenFog architecture will enable futuristic services and performance.

## Emergency / Disaster Response

Emergencies, like a fire or an active shooter, and disasters, like a flood or an earthquake, happen while people are in buildings at work, or while shopping. Though warning and surveillance systems have been widely applied to help the public be aware of such situations, existing warning systems have not been integrated into the buildings themselves.

This means that automatic actions that could save lives cannot occur without manual intervention, and information that could greatly aid the efforts of relief workers cannot be provided.

A Smart Building dynamically reacts to emergencies and disasters to protect its occupants and assist first responders, and it will need systems designed on the OpenFog architecture to do so.

Say there is a grease fire in the kitchen of a 6-story office building that gets out of control. Today's analog systems (if properly maintained) will have smoke detectors send electrical impulses down a dedicated fire alarm cable network to a central panel that triggers alarms throughout the entire building.

Occupants hear the alarms and, due to prior training or education, get up from their desks and hurry towards what they assume is the closest, safest, and most efficient exit. Some exits will draw clusters of people that must wait in line and slowly descend the stairs, while others will see no people at all. It may not be obvious if some exits are blocked by fire, smoke or debris.

Most modern buildings have a trigger to alert the local Fire Department automatically, but not all. The systems are designed as well as they could be, given the constraints, but are not very efficient due to their heavy dependence on all occupants knowing what to do,

and on all analog functions in their life safety networks operating as expected.

Now imagine the same scenario, but in a Smart Building that has fully embraced the OpenFog architecture. Smoke alarms or CCTV cameras (supported by fog-based video analytics) in the kitchen detect a fire spreading past normal levels. The discovery of the fire event could occur on local fog nodes in a matter of milliseconds.

The local fog node directs the kitchen's HVAC system to take action to minimize the chance of the fire spreading, and cuts off power to the circuits that could worsen the situation. It sends an alert to the rest of the fog nodes in the area estimated to be in danger instead of the entire building. The lights in those areas turn red, all screens (personal phones, laptops, TVs, etc.) in those areas send evacuation alerts, and alarm noises begin to sound.

The security system knows in the affected areas exactly how the people are distributed, and where the fire exits are. By pushing this live data through a pre-trained machine learning algorithm, the local fog node calculates the fastest exit path for each zone of people, and continues calculating it as they get up and move.

Cameras and sensors in stairwells and near exits ensure evacuation paths are safe, and reroute occupants if they are not. The screens and digital signs in the area all show visual exit paths based on this information, the overhead lights start strobing in the direction of the exit path so that people will follow, and alarm noises in each area are converted to live instructions and directional audio cues on which exits to use.

The Fire Department has already been alerted by telephone, email and/or SMS, and are ready and waiting outside before the last person has exited the building. The first responders' mobile devices are in touch with the building's fog nodes to direct them to the emergency. They have been advised what to be prepared for, with live video from

the drones and CCTV cameras stitched together to provide a full picture in virtual or augmented reality to help orchestrate the emergency response.

Ideally, these measures contain the fire and everybody that was calculated to be in danger was evacuated, while those that were safe were left alone. No one needed any special knowledge, as the fog systems directly targeted the evacuees and got the message across in a variety of different mediums.

Fog alleviates the need for multiple layers of networks and communications within a building and streamlines decisions and comms. The local nature of the building's fog network made this scenario faster and more reliable than any cloud or edge-based solution could be.

### Optimizing an Employee's Performance

A typical employee's day at work, looks like this:

- Leaves home and drives to the office
- Parks in the adjacent parking garage
- Passes through security by swiping an RFID badge to enter the building
- Depending on the day, navigates to an office, or a meeting, or an open desk
- Works until lunch, interfacing with numerous screens and coworkers
- Eats either on or off campus
- Works until the end of the day, again interfacing with numerous screens, offices and people
- Passes through security to exit the building
- Finds car in the parking garage
- Leaves the office and drives home

Legacy analog buildings were designed to provide a constant, comfortable atmosphere. Internal systems such as lighting, HVAC, or

telecommunications were built in silos on separate networks to do their jobs with minimal feedback or deviation. Employees are generally comfortable, their days occurring without incident.

But in an OpenFog architected world, each of these events presents an opportunity to provide an experience that enhances their day and makes more efficient use of their time. For example:

- Upon arrival, the parking garage fog node directs the occupant to the open spot closest to their desk or to their first meeting, depending on the day's schedule.
- CCTV has already noticed the occupant's car, with the occupant behind the wheel, and the WiFi has authenticated with his various connected devices – so a badge is unnecessary.
- The occupant's schedule accounts for all appointments, booked conference rooms, and those of every other occupant, to make sure each person knows where to find their next meeting, or is directed to a quiet place to focus.
- Lunch can either be ordered from an app and delivered, or current and predicted café wait times can drive decision-making. These can be paired with preset or suggested diet and exercise plans, perhaps based upon the building's tracking of each occupant's activity level.
- To help the occupants focus and avoid the post-lunch slowdown, the building minimizes the CO2 content in the air, adjusts the lighting to match Circadian rhythms and inserts low levels of white noise in open office areas to minimize cross-talk.
- The building recommends optimal times to leave, with the best paths to take home, based on traffic patterns and appointments placed in the calendar, such as picking the kids up from school.

- And of course, if any of these features don't fit the needs of a certain occupant, mobile devices and wall-mounted touchpads communicate with the building's fog nodes to customize all settings.

Each of these actions requires tight synchronization between real-time feedback loops and global optimization systems. There must be a wide distribution of fog-based devices that can receive gigabytes of data per minute. They need to parse through parameters such as space temperature requests, space occupancy triggers, a pressure loss from a door opening, or fan motor percentage, and then send out the correct actions. Security and privacy are paramount, with no fear of being hacked or having malware injected, to keep their occupants delighted with their space.

In addition, data gathered from disparate sources must be shared with numerous different systems instantaneously. For example, security CCTVs, local WiFi, and an occupant's internal personal profile all help the fog security system authenticate an occupant. The HVAC and lighting systems also use that information, to prepare the employee's space to suit their preferences before they arrive.

The horizontal spreading of massive amounts of data to enable real-time actions is where fog deployments are at their best. Without a fog deployment, the building owner is forced to manage dozens of different vendors and their cloud-based architectures, all competing for bandwidth and sharing data via APIs. The alternative is to not provide these services at all.

With fog, there can also be a central system keeping track of higher-level building KPIs, anticipating weather changes and coordinating with other nearby buildings or utilities to manage shared systems like parking, traffic, grid usage, and disaster response. Fog nodes in the building have the compute, storage and security requirements to satisfy these requirements, as well as the connectivity and coordination to feed the larger cloud-based systems with the data needed to inform their difficult decisions.

## Advanced Services

The following advanced services can be built on an OpenFog architecture:

- Dynamic energy conservation
- Virtual reality teleconferencing
- Augmented reality wayfinding
- Drones as security guards.

Predicting the building's energy profile throughout the day and dynamically adjusting to minimize the impact of each system, from lighting to HVAC to plug loads and others, can save building owners millions of dollars per year in operational costs. However, they require massive amounts of data and instantaneous decision-making.

Enabling occupants to call into a video conference and actually feel present in the room with the other callers will enhance productivity during those meetings. Providing visitors to a new building with an AR map to their next meeting, to a coworker's desk location, or to a nurse's office, will make them feel like they are in a truly "next generation" Smart Building. Designing autonomous drones that can patrol large factories from the air, and rapidly reach places a human cannot, will greatly improve an owner's confidence that their proprietary information and property is safe.

All of these services demand faster than 10 ms of latency for each user, with anywhere from 20 Mb/s to 1 Gb/s of throughput. If a cloud were governing the architecture, multiple instances of each service occurring on each floor of an office building would put incredible stress on the building's network, and would completely consume the building's entire Internet bandwidth.

However, if the data and computational resources serving each of these services is hosted in fog nodes within the building, they can dynamically request and share computing resources amongst each other to load-balance, which makes the entire scenario much more

feasible. The computing resource type could be selected to optimize the performance of each application, for example using combinations of CPUs, GPUs and FPGAs.

In this instance, the fog nodes receiving transmission from the drones will have enough intelligence to discern which images need to be stored and which can be discarded, greatly saving on bandwidth and storage capacity:

- The virtual conference room will mostly be defined by local resources, with only the information governing changes caused by users from a different location requiring internet connectivity, again saving on bandwidth.
- The wayfinding application will only pull data from the locally stored 3D model of the building, which is distributed in chunks amongst the different fog nodes so that no outside resources are required at all.
- Fog's low latency makes the AR application seem real. With an OpenFog architecture, these incredibly demanding services lose their ability to ruin the Smart Building experience for the rest of the occupants, as their resources are stored in the same place where their computing and networking demands are shared: locally.
- The autonomy of fog allows the building to continue these advanced features even if the cloud or internet backbone is down or overloaded
- Fog's interoperability insures a rich, dynamic supply base for the hardware, software and services that enable these innovations.

## Advantages of the Fog Computing Approach

The fog computing approach greatly reduces the latency, network bandwidth and availability constraints. The architecture for fog computing is based on eight pillars of elements, as identified by the OpenFog Consortium. In a Smart Building, the following pillars are addressed:




Figure 3. OpenFog Reference Architecture pillars mapped to the Smart Buildings Use Case.

**Security:** The Smart Building scenario is a physically distributed fog deployment, managing systems that could harm people if acted upon maliciously, or could compromise privacy. Any controls must be completely immune from access by a bad actor, and constant authentication and encryption is a must. Fog nodes are trustworthy enough to perform these functions.

**Scalability:** The fog system can adapt to the business’s needs as it relates to system cost and performance. When a new building system (e.g. badge readers, lighting, HVAC, or teleconferencing) and sensors and actuators are added, the solution must scale and not require a

completely new deployment. New, more demanding application mixes can be accommodated by expanding the capabilities of existing fog nodes, or adding more. As building occupancy increases, the building's fog infrastructure can grow with it.

**Open:** Openness is essential for the success of a ubiquitous fog computing ecosystem for IoT or 5G platforms and applications. Proprietary or single vendor solutions can result in limited supplier diversity, which can have a negative impact on system cost, quality and innovation. Fog systems, especially those based on an OpenFog-compliant architecture, can achieve seamless multi-vendor interoperability.

**Autonomy:** Autonomous data gathering, analysis, and actions are what makes a building smart. Fog deployments allow intelligence currently relegated to the cloud to operate at the edge, and enable a range of silo-busting experiences previously regarded as novelty. Since most key decisions are made locally with fog, the building's services continue without impairment even if the cloud is down, overloaded, or unreachable.

**Reliability/Availability/Serviceability:** The entire Smart Building system must be reliable, available, and serviceable, which includes orchestration of existing or new resources. As new energy efficiency models are trained for specific building types, they should be updated on or near edge devices without downtime to the system. Some smart building systems, for example elevator control, fire suppression, or active shooter mitigation are life critical. To support them, the fog infrastructure provides system availability better than 99.999%.

**Agility:** Smart Buildings must be dynamic, growing systems that constantly improve their operation. As new building sensors, actuators, and fog-based applications are added, the building's fog system seamlessly adapts to the new capabilities. These systems are too massive and heterogeneous to require hands-on maintenance or set-up for each piece of new equipment.

**Programmability:** Programming at the hardware level is often necessary for Smart Buildings to perform real-time fault detection and respond instantaneously. Many different stakeholders can program smart building fog systems, including the building’s architects, systems manufacturers, systems installers, landlords, maintenance staff, tenants, or individual occupants.

Fog Elements in Smart Buildings

As shown in Figure 4, fog elements in Smart Buildings span room, floor, building, and cloud nodes.




Figure 4. Fog elements in a smart building.

**Room Nodes**

Room fog nodes can govern a conference room, an office, an apartment or any similarly-structured closed space. These rooms can contain sensors monitoring temperature, humidity, air quality, occupancy, lighting, energy usage from plug loads and numerous other factors.

The room node will have enough intelligence to discern which data requires storage, based on which actions require real-time processing, which can be passed northbound to a floor fog node, or which can be shared east or west with another room node.

It will have enough analytical power to perform actions such as conserving plug load and HVAC usage when the room is unoccupied, to learn an occupant's preferred temperature and maintain the space at that comfort level, and to recognize intruders and sound an alarm. It can also have the sophisticated analytics ability to recognize gestures or voice commands requesting services such as casting a presentation from a laptop onto a screen, or conferencing in a colleague to a video meeting.

To keep these nodes cost-effective, especially considering how many will be required per building, when well-architected these should have just enough storage and computing to accomplish the aforementioned tasks. More complicated compute tasks should be left to the floor, building, and cloud fog nodes.

## **Floor Nodes**

Floor fog nodes connect to many of the same sensors and actuators as room fog nodes (those that are in open areas), and to room nodes themselves. They coordinate all of the room nodes for a given floor, exchange information with other floor nodes, and share with the building fog node any information required for optimizing the entire building, or for storage in the cloud. These enhanced capabilities require more advanced analytics and larger storage space than room fog nodes and maintain similar latency.

The well-architected floor node should provide enough storage to maintain status and history for each sensor, actuator and room node to support:

- A 20,000 square-foot open office area with 100 constantly moving people.
- Three to ten conference rooms that vary in occupancy status throughout the day.

The floor node's analytical power enables constant solving of complex optimization problems such as:

- Minimizing HVAC and lighting energy usage while maximizing occupant comfort,
- Load-balancing when half of the occupants are streaming 4K VR over WiFi and the other half are simply sending emails and browsing the Internet, and
- Providing an internal air-traffic control system for autonomous drones.

Each of these services must be able to act at the same time, and some of them (particularly the drones) must be reliable 99.999% of the time with some services (particularly AR / VR) needing less than 7ms of latency. This requires models trained at the building or cloud level pushed and stored on the local device, plus enough computing power to send out hundreds of messages and actions per second to maintain the Smart Building feeling.

Floor fog nodes may be closely associated with the building's Wi-Fi access infrastructure. They can perform fairly complex networking functions such as deep packet inspection, location based services, and continuous security scans for all devices connected to them.

Using fog's high performance, localized computing resources, the signal strengths, directions, and arrival times from each mobile device to the building's Wi-Fi can be analyzed, and location precision to the centimeter is possible. The storage at the floor level may act like a

local file server, web proxy or media cache. The computation in this fog node could be a shared resource, supplementing room nodes, low power sensors or mobile devices with serious compute power upon demand.

Though more complex than room nodes, fewer floor nodes are required, given that a single floor can govern upwards of twenty rooms at once. They have the same southbound communications requirements plus room node integration. They will likely only interface with other OpenFog nodes to the north, east and west.

### **Building Nodes**

Well-architected building fog nodes communicate with fog nodes below them in the hierarchy, other OpenFog nodes in nearby buildings and web services. They ingest data from room and floor fog nodes and take slower, more deliberate actions such as setting equipment schedules, optimizing overall load on the building's systems, and communicating with the cloud.

Building nodes also coordinate the operation of the nodes below them in the hierarchy – for example stitching together the video from many security cameras, or coordinating an entire floor of light fixtures. They can manage the load balancing and fault tolerance of nearby floor or room fog nodes, moving applications around as lower-level nodes fail or become overloaded.

Most of the Smart Building experiences that tenants will enjoy are governed by the lower level nodes. But to make the system truly next level they must all work in concert. Because the building node only communicates southbound with other fog nodes, and consumes orders of magnitude greater amounts of data, more modern and standardized protocols like RESTful APIs should be used.

Real-time actions are not as vital at this level in comparison to management and orchestration. However, security, capacity, and

reliability are more vital at this level. The building fog node can even host the building's block chain for determining which devices are trusted and how long they have been trusted for, along with a transaction history for each device.

The room and floor nodes can carry a subset of the distributed ledger for each of their local trusted devices, but the full block chain would prove too heavy for fog nodes with such storage constraints. With this feature, a building node can generate its own energy model and predict future performance, uncover mechanical design deficiencies or offer suppliers a template for how best to add future features.

Building fog nodes communicate in an east or west direction with other fog nodes serving other buildings in their area. They govern services such as Weather Prediction, Transportation, Emergency Response or Smart Grid. These common, interoperable OpenFog interfaces allow for the modular construction of truly Smart Cities, with buildings and their fog infrastructures being an integral component.

To properly service their customers, Dynamic Smart Grids, for example, will need real-time information from all of the building nodes in an area about their current and projected power demands. Building nodes, in turn, take feedback from these external fog nodes and direct it to the proper users, which might be a chiller plant that needs to ramp down, an occupant seeking traffic times home, or a data network requesting more bandwidth.

## **Cloud Nodes**

Cloud fog nodes provide the analytics power to train predictive models with huge, stored datasets, compare a Smart Building's performance to that of like Smart Buildings, and generate the wisdom learned from Petabytes of data from tens to hundreds of buildings. Computationally intensive procedures such as providing Proof-Of-Work for a building's block chain or training deep neural networks make the most economic sense when they occur in an energy efficient data center in the cloud.

Smart Buildings require trained models for visual surveillance systems, predictive temperature setting, building pre-cooling and peak demand shaving, amongst others. The training data for these models will come from the sensors attached to low level nodes such as room and floor, get fed and tagged through the building nodes and be stored in the cloud, driving storage space measured in tens of terabytes.

Standardized metrics such as Energy Star Rating do a decent job of approximating building performance, but the massive amounts of data Smart Buildings generate will allow for new and better metrics that drive electrical, water and carbon conservation. These new standards need to be agreed upon, but for now will be distributed through cloud-based models rather than top down from government programs. Though all buildings are unique, as the sheer scale of building models uploaded to the cloud increases, so does the likelihood of an energy profile match.

Finally, cloud nodes will provide the ability for a Smart Building to connect to other fog nodes in the area without opening itself up to malicious hacks. With the cloud being the single point of reference outside the building (along with the devices brought inside the building that will either be trusted on the block chain or partitioned off into a guest network) the attack surface becomes very small and manageable.

These other fog nodes will allow a Smart Building to integrate itself into a Smart City, connecting its occupants to live, customized data about their surroundings in brand new ways.

## Mapping to OpenFog Communications Architecture




Figure 5. OpenFog communication for Smart Building applications highlighted at the node level across cross-cutting concerns.

### Performance & Scale

Quality of Service (QoS) is critical for Smart Building subsystems. Emergency/disaster response systems, described above, must take priority over all other data communications, and fog's connectivity framework will ensure this is synchronized between all elements in a room, floor, and building. QoS also includes latency requirements for the Advanced Services, described above, like augmented reality wayfinding and indoor autonomous drones. These services will demand high data rates and ultra-low latency to perform at the level expected from a Smart Building.

Additionally, the OpenFog Architecture will scale across a portfolio of Smart Buildings. Owners will quickly be able to compare performance

of an operation or service in one building to the same operation or service across the world via their Cloud Node.

## **Network**

The physical and network layers of the fog-based Smart Building include a variety of protocols. Legacy systems are built on wired protocols like POE, BACnet, Modbus, OPC-UA, MQTT, and LONworks, while newer systems are integrating edge devices via wireless protocols like WiFi, Zigbee, Z-wave, BLE, 6LowPAN and others. With Fog nodes having the interoperability to pair with all of these different devices, standard network switches, IP-based protocols and Ethernet links will connect room nodes to floor nodes to building nodes and allow Ethernet to again be the main protocol to link across the entire building.

Backhaul connections to cloud nodes will require various WAN links, depending on location.

## **Security**

Smart Buildings will house many critical services that must be protected against cyber-attacks. Emergency/disaster response systems, described above, and a building owner or tenant's proprietary IP cannot be compromised. The OpenFog connectivity framework should manage permissions and data accessibility to different types of occupants, services and devices. This means that a guest for the day should not be able to adjust setpoints on the fire alarm system, nor should a new, untrusted Smart Lighting sensor be able to push data to secured areas of the on-site datacenter.

## **Manageability**

System management for these use cases are critical to ensure security, system health and ease of overall management (e.g., device provisioning, app lifecycle). There will need to be a trusted connectivity framework level in place just for this "management

plane,” beyond the “application plane” supporting the actual Smart Building Operations. Management operations will have to be highly automated to avoid unacceptable human labor content in the design, provisioning, installation and operation of the hundreds of fog nodes and tens of thousands of IoT endpoints that will serve a large building.

### **Data, Analytics & Control, IT Business and Cross Fog Applications**

Data analytics & control, IT business and cross-fog applications are key elements of the “application plane” for the fog computing system. For Smart Buildings, these are represented by the employee optimization services above.

To function effectively in a fog environment, these functions need to be linked together by a Connectivity Framework or Messaging Foundation into the overall application. All eight of the OpenFog pillars come into play in this cross-cutting function. The functionality for the overall system, on the application plane, is almost completely instantiated as multiple apps running on different compute nodes. It’s a fundamentally distributed series of apps working together closely from the edge to the core/backend compute nodes.

### Smart Buildings in Smart Cities

The Smart Building use case overlaps Smart Cities. The urban revolution cannot be better characterized than through this convergence.

The key to building safety, efficiency and intra-city communications is interoperability, which is fostered by fog. The evolution will feature smart cities using the data available from smart buildings and other infrastructure to use resources more efficiently and safely. There are significant changes in our future and fog computing is at the center of this evolution. The OpenFog RA defines an open data format,

interfaces and messaging standards designed to ensure this interoperability.




Figure 6. Smart Cities will feature intra-city communications and interoperability between smart buildings and smart infrastructure.

### Fast Track to OpenFog Testbeds

There is a large, growing ecosystem of device manufacturers and suppliers for smart lighting solutions, HVAC controls, video surveillance systems, and many of the other Smart Building capabilities described above. The OpenFog Consortium will create Testbeds for this ecosystem where the common interfaces developed for interoperability in fog will be standardized and certified. These unique spaces will foster cooperation among suppliers, and enable amazing applications to be developed using an architecture designed for low-latency, high-bandwidth, heterogeneous operations.

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 Reference Architecture

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](mailto: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
Ryan Gentry, Intel	Evan Birkhead, OpenFog Consortium
Chuck Byers, Cisco Systems	Judith Kelley, 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.*