

NIST Special Publication 500-325

Fog Computing Conceptual Model

*Recommendations of the National Institute of Standards and
Technology*

Michaela Iorga
Larry Feldman
Robert Barton
Michael J. Martin
Nedim Goren
Charif Mahmoudi

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.500-325>

NIST
**National Institute of
Standards and Technology**
U.S. Department of Commerce

NIST Special Publication 500-325

Fog Computing Conceptual Model

Recommendations of the National Institute of Standards and Technology

Michaela Iorga
Nedim Goren
*Computer Security Division
Information Technology Laboratory*

Larry Feldman
G2, Inc.

Robert Barton
Cisco

Michael J. Martin
IBM Canada, Ltd.

Charif Mahmoudi
*Advanced Network Technologies Division
Information Technology Laboratory*

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.500-325>

March 2018



U.S. Department of Commerce
Wilbur L. Ross, Jr., Secretary

National Institute of Standards and Technology
Walter Copan, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

National Institute of Standards and Technology Special Publication 500-325
Natl. Inst. Stand. Technol. Spec. Publ. 500-325, 14 pages (March 2018)
CODEN: NSPUE2

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.500-325>

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analysis to advance the development and productive use of information technology.

Abstract

Managing the data generated by Internet of Things (IoT) sensors and actuators is one of the biggest challenges faced when deploying an IoT system. Traditional cloud-based IoT systems are challenged by the large scale, heterogeneity, and high latency witnessed in some cloud ecosystems. One solution is to decentralize applications, management, and data analytics into the network itself using a distributed and federated compute model. This approach has become known as fog computing. This document presents the conceptual model of fog and mist computing and how they relate to cloud-based computing models for IoT. This document further characterizes important properties and aspects of fog computing, including service models, deployment strategies, and provides a baseline of what fog computing is, and how it may be used.

Keywords

cloud computing; cloudlet; edge computing; fluid computing; fog computing; fluid computing; Internet of Things (IoT); mist computing.

Acknowledgments

The authors would like to thank their colleagues and the experts in industry and government who contributed their thoughts to the creation and review of this document. Special thanks go to Chuck Byers of Cisco and Gill Hughes of Capgemini for their insightful suggestions, and to Jim St. Pierre, Matthew Scholl, David Ferraiolo, Lee Badger, John Messina and Eric Simmon of NIST, for their final technical and editorial comments on the content of this document.

Audience

The intended audience of this document is system planners, system architects, system engineers, system managers, program managers, technologists and networking specialists that consume or provide Internet of Things solutions leveraging cloud and/or fog computing services.

Table of Contents

1 Introduction 1

 1.1 Purpose and Scope 1

2 The Fog Computing Conceptual Model 2

 2.1 Fog Computing Definition 2

 2.2 Fog Node Definition 3

 2.3 Fog Computing Essential Characteristics 3

 2.4 Additional Characteristics Often Associated with Fog Computing 4

 2.5 Fog Node Attributes 4

 2.6 Fog Node Architectural Service Models 4

 2.7 Fog Node Deployment Models 5

 2.8 Mist Computing as Lightweight Fog Layer 6

 2.9 Mist Computing Definition 6

Annex A— Fog Computing vs. Edge Computing 7

Annex B— Acronyms 8

List of Figures

Figure 1 – Fog computing supporting a cloud-based ecosystem for smart end-devices. 2

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.500-325>

1 Introduction

Ubiquitous deployment of *smart*, interconnected devices¹ is estimated to reach 50 billion units by 2020². This exponential increase is fueled by the proliferation of mobile devices (e.g. mobile phones and tablets), *smart* sensors serving different vertical markets (e.g. *smart* power grids, autonomous transportation, industrial controls, *smart* cities³, wearables, etc), wireless sensors and actuators networks. New concepts and technologies are needed to manage this growing fleet of Internet of Things (IoT) devices.

1.1 Purpose and Scope

The acute need of the multitude of smart, end-user IoT devices and near-user edge devices to carry out, with minimal latency, a substantial amount of data processing and to collaborate in a distributed way, triggered technology advancements towards adaptive, decentralized computational paradigms that complement the centralized cloud computing model serving IoT networks.

Researchers, computer scientists, system and network engineers developed innovative solutions to fill the technological gaps. These solutions provide faster approaches that gain better situational awareness in a far more timely manner. Such solutions or computational paradigms are referred to as *fog computing*, *mist computing*, *cloudlets*⁴, or *edge computing*^{5,6}. Since no consensus exists on distinction among these concepts at the time this document was created, the authors considered it imperative to provide a definition that can be used by practitioners and researchers to facilitate meaningful conversations on the topic.

This document provides the conceptual model of *fog computing* and its subsidiary *mist computing*, and aims to place these concepts in relation to *cloud computing*⁷ and *edge computing*.

Additionally, the document introduces the notion of a *fog node* and the *nodes federation model* composed of both, distributed and centralized, often hierarchical clusters of fog nodes operating in harmony. This model is introduced as a building-block architectural approach for constructing, enhancing or expanding the *fog* and *mist computing* layers.

Furthermore, the document characterizes important aspects of *fog computing* and is intended to serve as a means for broad comparisons of fog computing capabilities, service models and deployment strategies, and to provide a baseline for discussion of what *fog computing* is and the way it may be used.

The capabilities, service types and deployment models form a simple taxonomy that is not intended to prescribe or constrain any particular method of deployment, service delivery, or business operation.

¹ https://en.wikipedia.org/wiki/Smart_device

² https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf

³ https://en.wikipedia.org/wiki/Smart_city

⁴ <https://en.wikipedia.org/wiki/Cloudlet>

⁵ https://en.wikipedia.org/wiki/Edge_computing; <https://opensource.com/article/17/9/what-edge-computing>

⁶ <https://www.openfogconsortium.org/10-areas-where-fog-extends-iot-5g-ecosystems/>

⁷ NIST Special Publication 800-145, *The NIST Definition of Cloud Computing*, <https://doi.org/10.6028/NIST.SP.800-145>

2 The Fog Computing Conceptual Model

2.1 Fog Computing Definition

Fog computing is a layered model for enabling ubiquitous access to a shared continuum of scalable computing resources. The model facilitates the deployment of distributed, latency-aware applications and services, and consists of *fog nodes*⁸ (physical or virtual), residing between *smart* end-devices and centralized (cloud) services. The *fog nodes* are context aware and support a common data management and communication system. They can be organized in clusters - either vertically (to support isolation), horizontally (to support federation), or relative to *fog nodes*' latency-distance to the *smart* end-devices. *Fog computing* minimizes the request-response time from/to supported applications, and provides, for the end-devices, local computing resources and, when needed, network connectivity to centralized services.

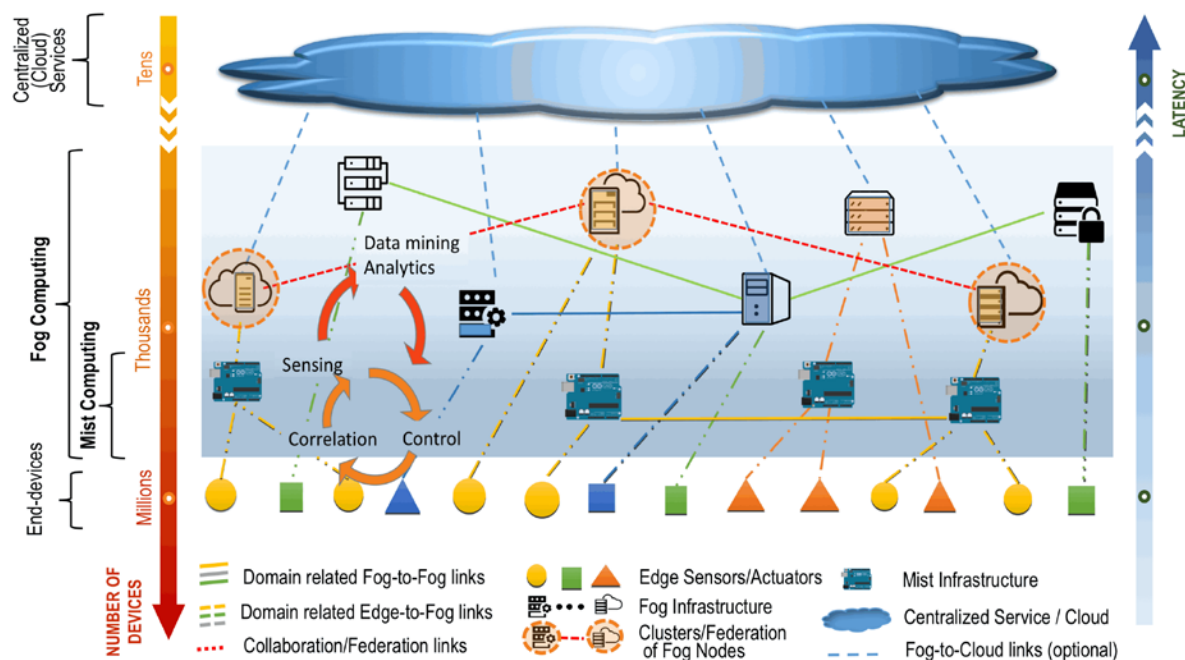


Figure 1 – Fog computing supporting a cloud-based ecosystem for smart end-devices.

Figure 1 above depicts fog computing in the broader context of a cloud-based ecosystem serving *smart* end-devices. *Fog computing* is not perceived as a mandatory layer for such ecosystems nor is the centralized (cloud) service perceived as being required for a *fog computing* layer to support the functionality of *smart* end-devices. Different usecase scenarios might have different architectures based on the optimal approach to supporting end-devices functionality. The choice

⁸ See Section 2.2 for the definition of a fog node.

of such representation is rooted in the intend of capturing a complex architecture that incorporates fog computing services.

2.2 Fog Node Definition

The *fog node* is the core component of the *fog computing* architecture. *Fog nodes* are either *physical* components (e.g. gateways, switches, routers, servers, etc.) or *virtual* components (e.g. virtualized switches, virtual machines, *cloudlets*⁹, etc.) that are tightly coupled with the *smart* end-devices or access networks, and provide computing resources to these devices. A *fog node* is aware of its geographical distribution and logical location within the context of its cluster. Additionally, *fog nodes* provide some form of data management and communication services between network's edge layer where end-devices reside, and the *fog computing* service or the centralized (cloud) computing resources, when needed. To deploy a given *fog computing* capability, *fog nodes* operate in centralized or decentralized manner and can be configured as stand-alone *fog nodes* that communicate among them to deliver the service or can be federated to form clusters that provide horizontal scalability over disperse geolocations, through mirroring or extension mechanisms.

2.3 Fog Computing Essential Characteristics

The following six characteristics are essential in distinguishing fog computing from other computing paradigms. However, a *smart* end-device or IoT user is not required to make use of all characteristics when consuming a *fog computing* service.

Contextual location awareness, and low latency. Fog computing offers the lowest-possible latency due to the fog nodes' awareness of their logical location in the context of the entire systems and of the latency costs for communicating with other nodes. The origins of *fog computing* can be traced to early proposals supporting endpoints with rich services at the edge of the network, including applications with low latency requirements. Because *fog nodes* are often co-located with the *smart* end-devices, analysis and response to data generated by these devices is much quicker than from a centralized cloud service or data center.

Geographical distribution. In sharp contrast to the more centralized cloud, the services and applications targeted by the *fog computing* demand widely, but geographically-identifiable, distributed deployments. For instance, the *fog computing* will play an active role in delivering high quality streaming services to moving vehicles, through proxies and access points geographically positioned along highways and tracks.

Heterogeneity. *Fog computing* supports collection and processing of data of different form factors acquired through multiple types of network communication capabilities.

Interoperability and federation. Seamless support of certain services (real-time streaming services is a good example) requires the cooperation of different providers. Hence, *fog computing* components must be able to interoperate, and services must be federated across domains.

Real-time interactions. *Fog computing* applications involve real-time interactions rather than batch processing.

⁹ <https://en.wikipedia.org/wiki/Cloudlet>

Scalability and agility of federated, fog-node clusters. *Fog computing* is *adaptive* in nature, at cluster or cluster-of-clusters level, supporting elastic compute, resource pooling, data-load changes, and network condition variations, to list a few of the supported *adaptive* functions.

2.4 Additional Characteristics Often Associated with Fog Computing

Predominance of wireless access. Although *fog computing* is used in wired environments, the large scale of wireless sensors in IoT demand distributed analytics and compute. For this reason, *fog computing* is very well suited to wireless IoT access networks.

Support for mobility. It is essential for many *fog computing* applications to communicate directly with mobile devices, and therefore support mobility techniques, such as the Locator/ID Separation Protocol (LISP)¹⁰, that decouple host identity from location identity, and require a distributed directory system.

2.5 Fog Node Attributes

To facilitate the deployment of a fog computing capability that exhibits the six essential characteristics described in Section 2.3, fog nodes need to support one or more of the following attributes:

Autonomy. *Fog nodes* can operate independently, making local decisions, at the node or cluster-of-nodes level.

Heterogeneity. *Fog nodes* come in different form factors, and can be deployed in a wide variety of environments.

Hierarchical clustering. *Fog nodes* support hierarchical structures, with different layers providing different subsets of service functions while working together as a continuum.

Manageability. *Fog nodes* are managed and orchestrated by complex systems that can perform most routine operations automatically.

Programmability. *Fog nodes* are inherently programmable at multiple levels, by multiple stakeholders - such as network operators, domain experts, equipment providers, or end users.

2.6 Fog Node Architectural Service Models

Fog computing, similar to the traditional cloud computing model, offers implementations of the architecture in multiple layers of the network's topology. Similar to the cloud computing service models defined in NIST Special Publication (SP) 800-145, the following types of service models can be implemented:

¹⁰ Locator/ID Separation Protocol (LISP): is a network architecture and set of protocols that implements a new semantic for IP addressing. LISP creates two namespaces and uses two IP addresses: Endpoint Identifiers (EIDs), which are assigned to end-hosts, and Routing Locators (RLOCs), which are assigned to devices (primarily routers) that make up the global routing system. See http://lisp.cisco.com/lisp_over.html for more information

Software as a Service (SaaS). The capability provided to the fog service customer is to use the fog provider's applications running on a cluster of federated fog nodes managed by the provider. This type of service is similar to the cloud computing Software as a Service (SaaS) and implies that the end-device or smart thing accesses the fog node's applications through a thin client interface or a program interface. The end-user does not manage or control the underlying fog node's infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS). The capability provided to the fog service customer is similar to the cloud computing Platform as a Service (PaaS) and allows deployment onto the platforms of federated fog nodes forming a cluster, of customer-created or acquired applications created using programming languages, libraries, services, and tools supported by the fog service provider. The fog service customer does not manage or control the underlying fog platform(s) and infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Infrastructure as a Service (IaaS). The capability provided to the fog service customer is to provision processing, storage, networks, and other fundamental computing resources leveraging the infrastructure of the fog nodes forming a federated cluster. Similar to cloud computing Infrastructure as a Service (IaaS) services, the customer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying infrastructure of the fog nodes cluster but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls) .

2.7 Fog Node Deployment Models

Since fog computing is identified and defined as an extension of the traditional cloud-based computing model, the following deployment models are also supported:

Private fog node. A *fog node* that is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units.) It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Community fog node. A *fog node* that is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations.) It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Public fog node. A *fog node* that is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the fog provider.

Hybrid fog node. A complex *fog node* that is a composition of two or more distinct fog nodes (private, community, or public) that remain unique entities, but are bound together by

standardized or proprietary technology that enables data and application portability (e.g., fog bursting for load balancing between these fog nodes.)

2.8 Mist Computing as Lightweight Fog Layer

Fog computing solutions are adopted by many industries, and efforts to develop distributed applications and analytics tools exist and continue to develop. The need for geographically dispersed, low-latency computational resources triggered the technological evolution of *fog computing* promoting development of more specialized, dedicated nodes that exhibit low computational resources. These nodes referred to as *mist nodes*, are perceived as *lightweight fog nodes*. These *mist nodes* that form the *mist computing* layer are placed even closer to the peripheral devices than the more powerful *fog nodes* they collaborate with, often sharing the same locality with the smart end-devices they service.

2.9 Mist Computing Definition

Mist computing is a lightweight and rudimentary form of *fog computing* that resides directly within the network fabric¹¹ at the edge of the network fabric, bringing the fog computing layer closer to the smart end-devices. *Mist computing* uses microcomputers and microcontrollers to feed into *fog computing* nodes and potentially onward towards the centralized (cloud) computing services.

The *mist computing* layer is not viewed as a mandatory layer of fog computing. When implemented, *mist nodes* can leverage the deployment models described in Section 2.5 and the service types described in Section 2.4.

¹¹ Network fabric is an industry term that describes a [network topology](#) in which components pass data to each other through interconnecting switches.

Annex A—Fog Computing vs. Edge Computing

For the purpose of this document, the *Edge computing* is the network layer encompassing the end-devices and their users, to provide, for example, local computing capability on a sensor, metering or some other devices that are network-accessible. This peripheral layer is also often referred to as IoT network.

Fog computing is often erroneously confused with edge computing, but there are key differences¹² between the two concepts. Fog computing runs applications in a multi-layer architecture that decouples and meshes the hardware and software functions, allowing for dynamic reconfigurations for different applications while performing intelligent computing and transmission services. Edge computing runs specific applications in a fixed logic location and provides a direct transmission service. Fog computing is hierarchical, where edge computing tends to be limited to a small number of peripheral devices. Moreover, in addition to computation, and networking, fog computing also addresses storage, control and data-processing acceleration.

¹² See OpenFog Consortium's whitepaper: <https://www.nebbiolo.tech/wp-content/uploads/whitepaper-fog-vs-edge.pdf>

Annex B—Acronyms

Selected acronyms and abbreviations used in this paper are defined below.

IaaS	Infrastructure as a Service
IoT	Internet of Things
PaaS	Platform as a Service
SaaS	Software as a Service
LISP	Locator/ID Separation Protocol