NIST Special Publication 800-191 (Draft)

The NIST Definition of Fog Computing

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

C O M P U T E R S E C U R I T Y

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

NIST Special Publication 800-191 (Draft)

The NIST Definition of Fog Computing

Michaela Iorga
Computer Security Division
Information Technology Laboratory

Larry Feldman
G2 Inc.

Robert Barton
Cisco

Michael J Martin
IBM Canada Ltd.

Nedim Goren
Computer Security Division
Information Technology Laboratory

Charif Mahmoudi
Advanced Network Technologies Division
Information Technology Laboratory

August 2017

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

National Institute of Standards and Technology
Kent Rochford, Acting NIST Director and Under Secretary of Commerce for Standards and Technology

Authority

This publication has been developed by NIST in accordance with its statutory responsibilities under the Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 *et seq.*, Public Law (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official. This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

National Institute of Standards and Technology Special Publication 800-191
Natl. Inst. Stand. Technol. Spec. Publ. 800-191, 13 pages (August 2017)
CODEN: NSPUE2

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 NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at http://csrc.nist.gov/publications.

Public comment period: *August 21, 2017* through *September 21, 2017*

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930
Email: SP800-191@nist.gov

All comments are subject to release under the Freedom of Information Act (FOIA).

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. ITL's responsibilities include the development of technical, physical, administrative, and management standards and guidelines for the cost-effective security and privacy of sensitive unclassified information in Federal computer systems. This Special Publication 800-series reports on ITL's research, guidance, and outreach efforts in computer security and its collaborative activities with industry, government, and academic organizations.

Abstract

Managing the data generated by Internet of Things (IoT) sensors 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 a formal definition 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 definition.

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 NIST Definition of Fog Computing 2
2.1 Fog Computing Definition 2
2.2 Fog Computing Characteristics 2
2.3 Fog Node Definition 3
2.4 Fog Node Architectural Service Types 4
2.5 Fog Node Deployment Models 4
3 Mist Computing as Lightweight Fog Layer 5
3.1 Mist Computing Definition 5

List of Appendices

Acronyms 7

List of Figures

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

1 Introduction

Ubiquitous deployment of smart, interconnected devices is estimated to reach 50 billion units by 2020[1]. 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 cities, 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 timelier manner. Such solutions or computational paradigms are referred to as *fog computing*, *mist computing*, *cloudlets*, or *edge computing*. Since no clear distinction among these concepts existed at the time the document was created, the authors considered it imperative to provide a formal definition that best matches the experts' views.

This document provides a formal definition of *fog computing* and its subsidiary *mist computing* concept, and aims to place these concepts in relation to *cloud computing*, *cloudlets* and *edge computing*.

Additionally, the document introduces the notion of a *fog node* and the *nodes federation model* composed of both, distributed and centralized 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 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://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf

2 The NIST Definition of Fog Computing

2.1 Fog Computing Definition

Fog computing is a horizontal, physical or virtual resource paradigm that resides between *smart* end-devices and traditional cloud or data centers. This paradigm supports vertically-isolated, latency-sensitive applications by providing ubiquitous, scalable, layered, federated, and distributed computing, storage, and network connectivity.

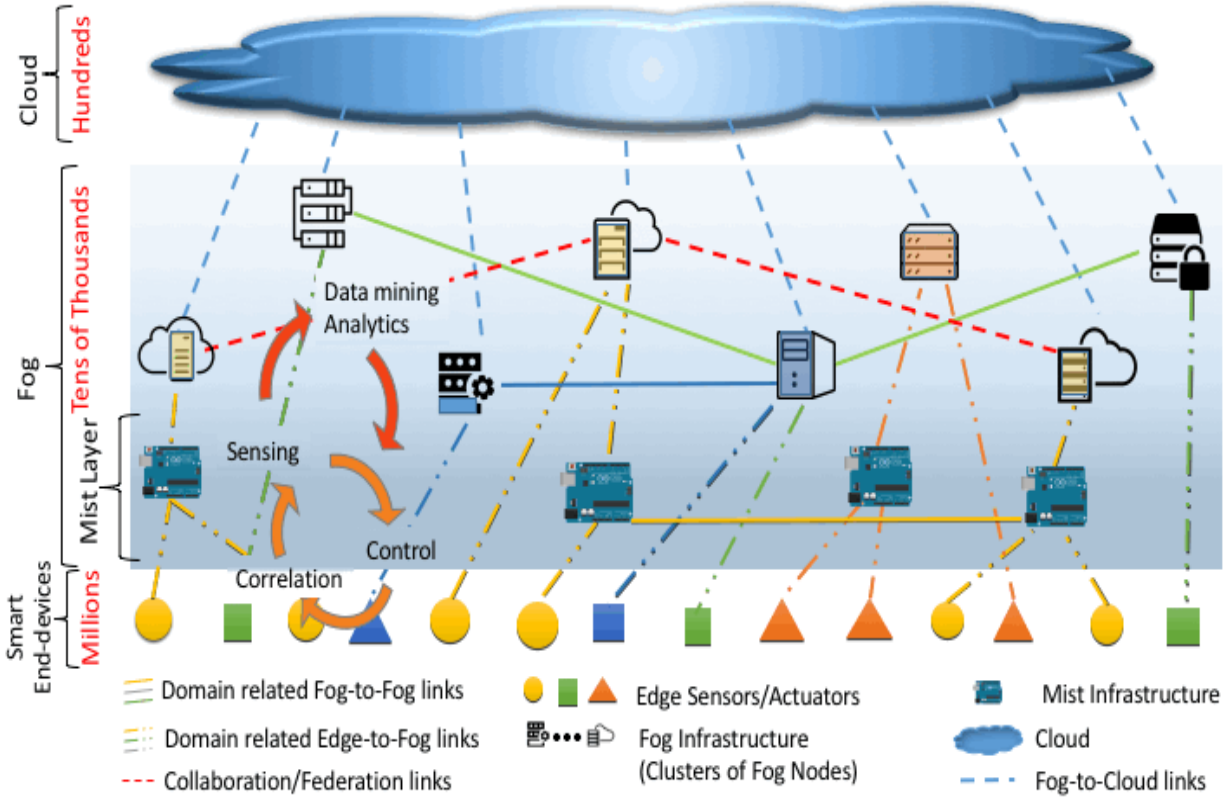

Figure 1 – Fog comuting supporting a cloud-based ecosystem for smart end-devices.

Figure 1 above depicts fog computing in the broader context of a cloud-based ecosystem serving smart end-devices. It is important to note that, in authors' view, fog computing is not perceived as a mandatory layer for such ecosystem.

2.2 Fog Computing Characteristics

Contextual location awareness, and low latency. The origins of the Fog can be traced to early proposals supporting endpoints with rich services at the edge of the network, including applications with low latency requirements (e.g. gaming, video streaming, and augmented

reality). Because Fog nodes tend to sit very close to the IoT endpoints, analysis and response to data generated by the endpoints is much quicker than from a centralized cloud.

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

Large-scale sensor networks to monitor the environment, and the Smart Grid are other examples of inherently distributed systems, requiring distributed computing and storage resources.

Very large number of nodes, as a consequence of the wide geo-distribution, as evidenced in sensor networks in general, and the Smart Grid in particular.

Support for mobility. It is essential for many Fog applications to communicate directly with mobile devices, and therefore support mobility techniques, such as the LISP protocol, that decouple host identity from location identity, and require a distributed directory system.

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

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 is very well suited to wireless IoT access networks.

Heterogeneity. Fog nodes come in different form factors, and will be deployed in a wide variety of environments, and the devices they collect data from may also vary in form factor and network communication capability.

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

Support for real-time analytics and interplay with the Cloud. The Fog is positioned to play a significant role in the ingestion and processing of the data close to the source as it is being produced. While Fog nodes provide localization, therefore enabling low latency and context awareness, the Cloud provides global centralization. Many applications require both Fog localization and Cloud globalization, particularly for analytics and Big Data. Fog is particularly well suited to real-time streaming analytics as opposed to historical, Big Data batch analytics that is normally carried out in a data center.

2.3 Fog Node Definition

Fog nodes are intermediatary compute elements of the smart end-devices access network that are situated between the Cloud and the smart end-devices. Fog nodes may be either *physical* or *virtual* elements and are tightly coupled with the smart end-devices or access networks. Fog nodes typically provide some form of data management and communication service between the peripheral layer where smart end-devices reside and the Cloud. Fog nodes, especially virtual ones,

also referred as *cloudlets*, can be federated to provide horizontal expansion of the functionality over disperse geolocations.

2.4 Fog Node Architectural Service Types

Fog computing is an extension of the traditional cloud-based computing model where implementations of the architecture can reside in multiple layers of a network's topology. Similar to cloud, the following types of service models can be implemented:

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 access 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 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).

2.5 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.)

3 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 disbursed, 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 and users than the more powerful fog nodes they collaborate with, often sharing same locality with the smart end-devices they service.

3.1 Mist Computing Definition

Mist computing is a lightweight and rudimentary form of computing power that resides directly within the network fabric[2] at the edge of the network fabric, the fog layer closest to the smart end-devices, using microcomputers and microcontrollers to feed into fog computing nodes and potentially onward towards the cloud computing services.

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

[2] 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* is the network layer encompassing the smart 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 also is often erroneously called edge computing, but there are key differences. Fog works with the cloud, whereas edge is defined by the exclusion of cloud and fog. Fog is hierarchical, where edge tends to be limited to a small number of peripheral layers. Moreover, in addition to computation, fog also addresses networking, storage, control and data-processing acceleration.

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